UC Santa Barbara
UC Santa Barbara Electronic Theses and Dissertations

Title
Essays in Applied Economics

Permalink
https://escholarship.org/uc/item/64x1q4qg

Author
Uhrig, Richard

Publication Date
2022

Peer reviewed|Thesis/dissertation

eScholarship.org Powered by the California Digital Library
University of California

University of California

Santa Barbara

Essays in Applied Economics

A dissertation submitted in partial satisfaction

of the requirements for the degree

Doctor of Philosophy

in

Economics

by

Richard Anton Uhrig

Committee in charge:

Professor Kelly Bedard, Chair
Professor Javier Birchenall
Professor Peter Rupert

June 2022

The Dissertation of Richard Anton Uhrig is approved.

Professor Javier Birchenall

Professor Peter Rupert

Professor Kelly Bedard, Committee Chair

June 2022

Essays in Applied Economics

Copyright © 2022

by

Richard Anton Uhrig

Acknowledgements

First and foremost, I would like to thank my entire dissertation committee for all of their advice, help, support, and guidance during my time at UCSB. I could not have asked for a better committee than Kelly Bedard, Javier Birchenall, and Peter Rupert. Each of you has helped me grow as a scholar here at UCSB, and I very much appreciate all the time and effort you all have spent to read my papers, watch me present, and provide endless feedback. Thank you.

Second, I would like to thank the rest of the faculty, administrative staff, and other graduate students with whom I have interacted over the past six years. I am grateful for our time together in classes, meetings, working groups, seminars, and various informal discussions about research, classes, teaching, life, and everything in between. I especially want to mention Mark Patterson for helping me navigate various administrata so I could focus on what I needed to. There are so many students, professors, and lecturers who have helped me in countless ways, and I appreciate all of you. Special recognition is deserved by friends and colleagues Ryan Sherrard, Molly Schwarz, Antoine Deeb, Jeffrey Cross, Maria Kogelnik, and Matthew Fitzgerald.

Third, I would like to thank my family and friends for all their support. I am especially thankful for my parents Cathy and Rick, my siblings Michael and Andrew, and all the friends from Santa Barbara and before that have had my back throughout this journey.

Last but not least, a special thanks to my partner Melissa Hsu, who has shared the vast majority of this experience with me and deserves an immense amount of credit for her kindness, patience, love, and encouragement.

Curriculum Vitæ

Richard Anton Uhrig

Education

2022	Ph.D. in Economics, University of California, Santa Barbara.
2017	M.A. in Economics, University of California, Santa Barbara.
2016	B.A. in Economics and Mathematics, College of William and Mary

Publications

Do Fans Impact Sports Outcomes? A COVID-19 Natural Experiment, with Jeff Cross. (*Forthcoming, 2022* **Journal of Sports Economics**)

Multiplicities: Adding a Vertex to a Graph, with Kenji Toyonaga and Charles Johnson. (*Vol 192, 2017* **Applied and Computational Matrix Analysis. Spring Proceedings in Mathematics and Statistics**)

Working Papers

The Effect of State-Level Rate Bill Abolition on School Attendance in the 19th Century United States

Honors and Awards

2022	UCSB Certificate in College and University Teaching
2020	UCSB Research Quarter Fellowship, Spring
2017	UCSB Williamson Fellowship

UCSB Teaching Experience

Teaching Associate, ECON 134A: Financial Management; Winter 2019, Winter 2020

Teaching Assistant, ECON 204A: Macro I (Ph.D Level); Fall 2019, Fall 2020

Teaching Assistant, ECON 2: Principles of Macroeconomics; Winter 2018, Spring 2018, Spring 2022

Teaching Assistant, ECON 5: Statistics for Economists; Fall 2021

Teaching Assistant, ECON 101: Intermediate Macroeconomic Theory; Spring 2017

Teaching Assistant, ECON 114B: Economic Development; Winter 2020

Teaching Assistant, ECON 134A: Financial Management; Fall 2018, Summer 2019, Summer 2020

Teaching Assistant, ECON 181: International Finance; Winter 2021, Spring 2021, Winter 2022

Teaching Assistant, PSTAT 109: Statistics for Economics; Winter 2017, Fall 2017

Conferences and Seminars

2021	Applied Micro Lunch, UCSB
2020	Applied Micro Lunch, UCSB
2020	Economic History Coffee Hour, UC Davis

Programming Skills

Stata, Python, LaTeX

Languages

English (Native), German (Intermediate)

Permissions and Attributions

1. The content of Chapter 3 and Appendix A.3 is the result of a collaboration with Jeffrey Cross of Hamilton University and is a previous version of an article that has previously appeared in the Journal of Sports Economics (Cross, J., Uhrig, R. (2022). Do Fans Impact Sports Outcomes? A COVID-19 Natural Experiment. Journal of Sports Economics. https://doi.org/10.1177/15270025221100204). It is shown here with the permission of the Journal of Sports Economics: `https://journals.sagepub.com/doi/full/10.1177/15270025221100204`.

Abstract

Essays in Applied Economics

by

Richard Anton Uhrig

This dissertation contains three chapters in applied economics. Chapter 1 discusses the abolition fo rate bills for public schools in the United States during the 19th century and the resulting impact on primary school attendance. Until the late 19th century, families in some municipalities paid small user fees, called rate bills, for their children to attend public schools. Urban school districts gradually repealed these fees and funded public education through local taxes. States eventually abolished rate bills, forcing rural areas to provide public education without tuition requirements. Using United States Census data and a staggered adoption difference-in-differences approach, I show that state-level rate bill abolition increased rural primary school attendance by 7.2 percentage points. These results suggest that small costs can be an obstacle to school attendance and inhibit the diffusion of education.

Chapter 2 discusses the effect of urbanicity on fertility. Fertility rates are significantly lower in urban areas than in rural areas. Using data from the 2016 Canadian Census, I exploit the exogenous placement of refugees to decompose the direct (causal) and indirect (compositional) effects of urbanicity on fertility. I find that these differences persist for refugee women, suggesting that causal factors are primarily responsible for rural-urban differences in fertility.

Chapter 3 discusses the effect of fans on home field advantage in European soccer. We exploit exogenous variation in the level of fan attendance driven by COVID-19 mitigation policies and find that the home field advantage, as measured by home minus away (expected) goals, is reduced by more than 50% across the English Premier League, German Bundesliga, Italian Serie A, and Spanish La Liga. This leads to a decrease in probability for a home win, indicating that these goals are pivotal with respect to match outcomes. These results are robust to controlling for factors potentially correlated with the no-fans policies, such as changes in weather and COVID-19 prevalence by region.

Contents

Chapter 1

The Effect of State-Level Rate Bill Abolition on School Attendance in the 19th Century United States

1.1 Introduction

During the late 19th century, the United States became one of the most educated societies in the world. 97.1% of all children between the ages of 5 and 14 were enrolled in primary school (Lindert, 2004). A large fraction of these children attended public school: 89.6% of children aged 5-14 were enrolled in public schools by 1910, compared to only 54.6% enrolled in public schools in 1830 (Lindert, 2004). Over this time period, public schools became free to attend through the abolition of tuition payments and other fees, shifting the burden of education funding from students and their families to all of society. This shift predates compulsory schooling laws (CSL's), which eventually made primary school mandatory to attend. This paper analyzes the impact of state-level rate bill abolition laws, which prevented public schools from charging tuition, on attendance in rural areas. Although these 19th century rate bills were small relative to income, I find that their abolition led to increased attendance in rural areas, implying that even low levels of tuition can hinder educational attainment.

There is limited evidence on the effect of rate bills on attendance in the United States. The 1826 imposition of a rate bill between 25 cents and 2 dollars per academic quarter in New York City led to a 13% drop in attendance (Cubberley, 1919). For those more advanced courses and higher associated tuition, the drop in attendance was above 90%. Go (2015) studies the effect of rate bill abolition in Connecticut and finds that larger initial rate bills, and therefore larger decreases at abolition, led to larger increases in attendance when the state banned tuition requirements in 1868. However, this is the first paper to investigate the abolition of rate bills across multiple states using the exogenous implementation of free public schools in rural areas.

Using data from the United States Decennial Censuses of 1850 through 1880, I investigate the effect of state-level rate bill abolition laws on attendance in rural areas using a staggered adoption difference-in-differences framework. These policies were effectively imposed upon rural areas, via state legislatures or referenda, by urban areas, which had already repealed rate bills prior to these state-level bans. I find that the abolition of rate bills increased rural attendance by 7.2 percentage points for individuals below the age of seven when the laws were passed. These results imply an increased average educational attainment of 0.72 years for children educated in rural areas,[1] approximately equal to the effect of compulsory schooling laws estimated in Black *et al.* (2008)[2] and approximately

[1]This number is calculated based on a 7.2 percentage point increase in attendance over ten years of childhood.

[2]However, these results for CSL's are statistically insignificant when clustering at the state level,

21% of the increased educational attainment driven by the expansion and diffusion of the American high school in the early 1900's (Goldin, 1998). The positive effects on attendance persist for individuals observed above the age of nine, indicating that the abolition of rate bills induced children to continue attending through late primary school and thus provided the foundation for the "high-school revolution" documented by Goldin (1998). These results suggest that the presence of tuition for post-secondary schools in the United States and primary and secondary schools in some developing countries, even at low levels, may depress attendance and educational attainment.

Increased attendance in response to rate bill abolition is somewhat surprising for a few reasons. Rate bills were not expensive, even for the time period. In New York state, average tuition for the full 18-week academic term in 1841-1842 was only 0.3% the annual wages of a non-farm worker (Go & Lindert, 2010). Based on the authors calculations, rate bills in the relatively expensive American South for an entire academic year were only 2.5% of annual wages for farmers.[3] While the marginal student's family presumably earned less than the averages used here, these statistics demonstrates the small magnitude of rate bills at this time. Additionally, tuition and user fees were not the only costs associated with primary schools at this time. Child labor laws were extremely rare in the 19th century, and some children worked in both rural and urban areas.[4] Therefore, non-attendance may be driven by the opportunity cost of children working rather than the pecuniary cost of rate bills. The elimination of rate bills only altered one aspect of the decision for marginal school attendees, leaving opportunity costs unaffected, but still led to increased attendance. Lastly, public schools charging rate bills were not

driven by large standard errors. In contrast, my results are robust to clustering at the state level and using the wild bootstrap to account for a low number of (effective) clusters.

[3]Data on rate bill magnitudes are from Go & Lindert (2010). Data on wages are from Census Bureau (1975).

[4]Limited data in the 1860-1880 Censuses suggest that between 4% and 6% of primary-school-age children worked overall, with slightly more working in rural areas than urban areas.

the only options for children seeking primary education. Some states and municipalities also provided pauper schools free of charge to those individuals on lower rungs of the socioeconomic ladder, for whom rate bills at public schools represented a higher fraction of income. For the wealthy, parents may choose to send children to private schools instead of tuition-charging public schools. In fact, "[rate bill] abolition was accompanied generally by a decreased attendance at private schools" (Cubberley, 1919), implying some level of substitution between these two alternatives.[5] Other primary school options meant that rate bills would have only inhibited a subset of individuals considering public schools. However, the removal of these charges still induced attendance among marginal students that could not attend pauper schools or afford private schools. Thus, the striking result here is that even small fees were large barriers to school attendance.

These results are economically and historically significant in the broader development of the United States throughout the 19th and 20th centuries. Relevant estimates for returns to education range from 3.8% (Goldin & Katz, 2000) to 10.6% (Duflo, 2001) per year. Therefore, the returns generated by an additional 0.72 years of educational attainment suggest increased income of between 2.7% and 7.6% for the rural populations induced to attend school by the abolition of rate bills. Thus, this shift in school funding contributed to the economic development of the United States in a significant way.

These policies could also induce structural change in the United States through three channels. First, high-education workers are more suited for non-agricultural types of work (Goldin & Katz, 2008). Second, increased productivity of the agricultural sector allows for "surplus" workers to shift into other sectors of the economy (Cao & Birchenall, 2013). Third, increased education is associated with higher innovation and research

[5]Specifically, this suggests that wealthier families were explicitly choosing to send children to private schools because they would be required to pay tuition for either option.

(Goldin & Katz, 2000; Toivanen & Väänänen, 2016), thus encouraging shifts towards more technology-intensive areas like industry.

Rate bill abolition was a major step in the shift towards accessibility of education for the entire population. Free public schools in the 19th century were sometimes referred to as "common schools" because they were available to all children that wished to attend (Cubberley, 1919; Goldin & Katz, 2008). Additionally, as is mentioned above, rate bill abolition sets the stage for the gains experienced in the first half of the 20th century as a result of the high school revolution (Goldin, 1998). This contrasts the earlier iterations of secondary schooling in the United States and abroad; secondary education was primarily used as preparation for college and university, as opposed to granting terminal degrees in their own right. By contributing to the shift towards universal primary schooling in rural areas, rate bill abolition may have been important in providing the foundation upon which the United States became the most educated country in the world in the late 19th and early 20th centuries.

This paper also contributes to the literature on discontinuities and non-linearities for prices around zero. This is relevant to the discussion of post-secondary education today, where community colleges charge very low levels of tuition. Denning (2017) shows that post-secondary attendance is more sensitive to decreased tuition when the price is already close to zero. Similarly, significant research in health economics (Harris *et al.*, 1990; Choudhry *et al.*, 2011; Baicker *et al.*, 2015; Gross *et al.*, 2020) show that small co-pays significantly reduce the probability a prescription is filled, even when the benefits of medication greatly outweigh the cost to the customer. In both cases, even low prices seem to matter to potential consumers of education or healthcare. This sentiment is echoed in my results, where the elimination of small rate bills leads to large attendance increases.

The rest of the paper proceeds as follows: Section 2 provides background on the education system in the United States during the the 19th century. Section 3 introduces the data utilized in this paper. Section 4 describes the staggered adoption difference-in-differences empirical strategy and the wild bootstrap. Section 5 highlights the key findings and provides various robustness checks. Section 6 discusses the economic and historical significance of rate bill abolition and the resulting increase in attendance. Section 7 concludes.

1.2 Background

Rate bills were first implemented by school districts as education systems expanded in the 18th and 19th centuries as an alternative to increased government funding, thus sharing the cost of public school provision between the the local population writ large and the specific families of students utilizing the service.[6] The magnitude of rate bills varied across states and school districts but was generally small relative to incomes. According to Go & Lindert (2010), rate bills comprised over 75% of public school funding in the state of New York in 1825. In the case of New York City, for which Cubberley (1919) provides the most detail, rate bills were charged based on classes; more advanced courses like astronomy and bookkeeping required much higher rate bills than elementary classes on alphabet and spelling. The largest of these fees was \$2, or the equivalent of approximately \$52 in 2019. Go & Lindert (2010) estimates that tuition for 18 weeks of schooling in 1841-1842 was approximately 0.3% of annual wages for the average non-farm worker in New

[6]"Starting in the 17th century, more and more localities [in the United States] developed their own school districts. Their funds came mainly from local property taxation, but also from tuition, donations, and occasional help from state land-sale revenues" (Lindert, 2004).

York state.[7] In that same school year, rate bills accounted for 41% of school revenues, with the rest coming from public funding.

Over time, the general population thought of public provision of education as more important to society. As these attitudes shifted, school districts reduced rate bill requirements to better promote school attendance. By 1850, rate bills provided only around 35% of public school funding in New York state, a significant decrease relative to 1825. Tuition and fees were primarily replaced by local taxes and state subsidies for public schools, shifting the cost from those families utilizing public education to all of society. This change towards fully-funded schools rather than partial tuition began primarily in cities.[8] Urban areas typically repealed rate bills of their own volition in an attempt to increase school attendance and educational attainment, just as New York City had done in the 1830's. Rural areas, on the other hand, continued to charge tuition and vote against rate bill abolition until state legislatures and public referenda forced their hand.

Indeed, when states began considering the abolition of rate bills throughout all school districts, urban areas were the primary drivers of this change.[9] In a 1849 referendum by the state of New York, the people voted to publicly fund schools. This referendum was reaffirmed by a second public vote the next year. In both cases, urban areas voted for rate bill abolition and rural areas voted for the status quo.[10] Despite these referenda,

[7] Go & Lindert (2010) find that it took 0.16 weeks of non-farm wage work to pay the rate bills for 18 weeks of school, which is 0.3% of a 52-week year. While a large fraction of the population, especially in rural areas, was involved in agriculture, non-farm wage data is more available for this time period.

[8] At least eleven cities in New York provided free schools before the state abolished rate bills in 1868: New York City, Buffalo, Hudson, Rochester, Brooklyn, Williamsburg, Syracuse, Troy, Auburn, Oswego, and Utica. Other major cities such as Baltimore, Charleston, Mobile, New Orleans, Louisville, Cincinnati, Chicago, and Detroit also provided free schools at least 25 years prior to state-level rate bill abolition (Cubberley, 1919).

[9] "Cities demanded educational progress, and were determined to have it, regardless of cost" (Cubberley, 1919).

[10] "[The cities] would not tolerate the rate bill [anywhere in the state], and, despite their larger property interests, they favored tax-supported schools" (Cubberley, 1919).

rural areas continued to charge tuition. The legislature for the state of New York finally abolished rate bills almost two decades later for those remaining communities that had continued to charge them. These rural-urban differences were typical in other states as well, as one might expect given that many urban areas in the rest of the country had repealed rate bills prior to state-level abolition; rural areas clearly had no qualms about rate bills, as they continued to charge tuition and fees well after cities provided primary schooling free of charge.

1.3 Data

Information on rate bill abolition laws is collected from Go (2009), which draws on many different sources, including Cubberley (1919), Goldin & Katz (2008), Swift (1911), and Mead (1918). Table 1 shows the dates of passage for rate bill abolition as well as first instances of state-level compulsory schooling laws, which are drawn from Goldin & Katz (2008). Figure 1 displays a map of the United States, highlighting those states used in various specifications. For the primary analysis, I include all 27 states for which rate bill abolition dates are available.[11]

I utilize the 1850-1880 Full Count Censuses from the Integrated Public Use Microdata Series (Ruggles *et al.*, 2019). Key variables in the data include information on an individuals attendance status, age, race, gender, various measures of socio-economic status, urbanicity,[12] and state of birth. I restrict my dataset to native-born whites between the ages of 5 and 14 in rural areas.[13] Table 2 shows relevant summary statistics

[11]Restrictions for specific robustness checks are discussed in the appropriate sections.

[12]Individuals are considered to be rural if there are fewer than 2,500 people in the municipality. Approximately 75% of the sample is considered rural by this metric.

[13]Children of other races or nationalities may have been subject to discrimination that prevented

for each of the four groups of states in my dataset: those that abolish rate bills prior to 1850, between 1850 and 1860, between 1860 and 1870, and between 1870 and 1880.

In each Census year utilized here, the Census asked some variation of "Was the person at school within the last year?" for every person in the household.[14] Figure 2 plots the attendance rates by age for individuals between the ages of 5 and 18, broken down across dimensions of gender and urbanicity. Men and women attended school at similar rates, regardless of where they lived. Attendance in urban areas was higher than rural areas, likely due to children working rather than attending school (Cubberley, 1919).In any case, rural children typically lagged behind urban ones in primary school attendance, possibly driven by the presence of rate bills,[15] child labor, and/or differences in returns to education.[16]

1.4 Empirical Strategy

1.4.1 Identification

In order for the effect for rate bill abolition to be identified, I assume that the legislation and referenda that brought about these changes were not accompanied by other education reform policies. It is necessary that the only aspect changing with

attendance regardless of any user fees charged. I choose ages 5 through 14 to focus on primary-school-age children. Urban areas are used as a placebo test since rate bill abolition was not binding to areas that already provided free public schools.

[14]The question included above was the question in 1850. The following questions were used in 1860, 1870, and 1880: "Did the person attend school within the last year?" "Did the person attend school within the last year?" and "Had the person attended school within the past year?" I treat these four as equivalent questions about one's school attendance.

[15]As is mentioned above, public schools were typically free in urban centers even prior to state-level abolition.

[16]Limited data in the 1860-1880 Censuses suggest that between 4% and 6% of primary-age children

treatment is the existence of rate bills, rather than the availability of education or the returns to education.[17] If this assumption were violated, then one would expect such legislation to affect attendance in urban areas as well as rural areas. Therefore, I use urban areas as a placebo test to check whether or not rate bill abolition laws led to changes in attendance in areas where rate bills had already been repealed. While I find strong, positive, statistically significant effects of rate bill abolition on attendance in rural areas, I estimate a null effect in urban areas. This provides some evidence reinforcing the assumption about other attendance-promoting policies; such policies would presumably affect urban areas as well, whereas rate bill abolition would only affect those places that had charged rate bills immediately prior to state-level abolition.

I also assume that individuals above the age of ten when these policies are implemented are unaffected by rate bill abolition. Presumably, the educational trajectory of an older individual is unlikely to be affected by minor changes in pecuniary costs, so the abolition of rate bills for the remaining years does not affect their attendance status going forward. This is consistent with evidence that attendance did not increase for children at older age groups unless these policies were implemented when they themselves were younger, shown in Table A.1 and A.2. If this assumption is violated, estimated coefficients could therefore be interpreted not as the true effect of rate bill abolition but instead as the effect of rate bill abolition on young children less the effect of rate bill abolition on older children, thus attenuating the results towards zero. My results are robust to various upper bounds on ages that would be affected by these policies.

Similarly, I assume that rural areas continued to charge rate bills until state-level policies abolished their use. Given the well-documented repeal of rate bills in urban areas

worked overall, with slightly more working in rural areas than urban areas.

[17]The availability of and returns to education almost surely changed over this time, but I only assume that these changes were orthogonal to treatment status.

and no corresponding evidence in rural areas, as well as the voting patterns in rural areas against rate bill abolition, I believe this to approximate the truth. If this assumption were violated, then estimated coefficients understate the true importance of rate bills as a barrier to attendance, since abolition would not be binding for the entire group and instead only affect that subset of rural areas that were forced to repeal rate bills by state-level policies.

1.4.2 Treatment

Treatment is assigned based on a child's age when rate bills were abolished in their state. For the primary specification, $Treatment = 1$ for individuals that were younger than seven years old when the rate bills were abolished, $PartialTreatment = 1$ for individuals that were between the ages of seven and ten years old inclusive when the rate bills were abolished, and individuals older than ten years old when these laws are passed are given both $Treatment = 0$ and $PartialTreatment = 0$. Of course, all individuals are considered untreated prior to rate bill abolition.

Consider Connecticut, which passed a rate bill abolition law in 1868. All individuals observed in the 1850 and 1860 Censuses are considered untreated. Children born in 1856 or 1857 are untreated in the 1870 Census, since they were likely unaffected by these policies that were implemented at ages 12 and 11, respectively. Individuals born between 1858 and 1861 are considered partially treated, since they presumably started school prior to rate bill abolition but the law took effect during their prime school-going years. Children born after 1861 are considered fully treated, whether they are observed in 1870 or 1880, because the law was passed before they turned seven.

Over the four Decennial Censuses used here, most individuals fall into either the "fully-treated" or "fully-untreated" categories. Only 8.4% of individuals are considered partially treated; 66.1% of children in my sample are fully treated, and 25.6% of individuals are fully untreated. Therefore, this analysis is primarily a comparison between individuals on either extreme that either started school after these rate bill abolition policies took effect or attended school at a time such that they did not especially benefit from the abolition of tuition payments. It is therefore unsurprising that results are not sensitive to alternative age cutoffs for full and partial treatment, as the fully-treated and untreated groups comprise the vast majority of the data.

1.4.3 Estimating Equation

I employ the following staggered adoption difference-in-differences specification with two-way fixed effects:

$$Y_{aist} = \alpha + \beta_1 Treatment_{ast} + \beta_2 PartialTreatment_{ast} + \gamma_{as} + \delta_{at} + \varepsilon_{aist} \tag{1.1}$$

where the outcome variable Y_{aist} is the attendance status of an individual i observed at age a in state s and Census year t, γ_{as} is a state-by-age fixed effect, δ_{at} is a year-by-age fixed effect, and ε_{aist} is the error term. Year-by-age fixed effects improve upon the combination of age and year fixed effects by allowing attendance differences between ages to vary across the four Census years considered here.[18] Similarly, state-by-age fixed effects improve upon state fixed effects because they account for differences in attendance

[18]My findings are robust to including age and year fixed effects instead of year-by-age fixed effects.

rates by age for each state, rather than simply accounting for the overall attendance level with no regard for how attendance differs by age.[19]

Errors are clustered at the state level (thus allowing for correlation in the error terms between individuals in the same state across all Census years) due to serial correlation in the treatment variable, as is prescribed by Bertrand *et al.* (2004). Therefore, I obtain my critical values from the t-distribution with $G-1=26$ degrees of freedom rather than from the normal distribution to account for the low number of clusters. Additionally, I utilize the wild bootstrap with Rademacher weights as prescribed in Cameron *et al.* (2008) to construct a 95% confidence interval for the effect of state-level rate bill abolition on education.[20]

1.5 Findings

1.5.1 Main Results

Column 1 of Table 3 shows the results of my primary specification. I find that rural children at younger ages when rate bill abolition laws come into effect are 7.2 percentage points more likely to attend school than the control group. The estimated coefficient is statistically significant at the 1% level when critical values are taken from the t-distribution with $G-1=26$ degrees of freedom to account for the low number of clusters. Confidence intervals from the wild bootstrap are shown in brackets below the standard errors in Table 3. The estimated coefficient remains significant at the 5% level using this more conservative methodology. The results presented here suggest that small

[19] My findings are robust to including state fixed effects instead of state-by-age fixed effects.

[20] The potential pitfalls of a low number of effective clusters have been well documented by Carter

rate bills were a large obstacle to primary school attendance and that their removal led to increased educational attainment in rural areas of the United States in the late 19th century.

Column 2 of Table 3 shows the results of Equation 1 estimated on urban areas, which provide a useful placebo test. Note that rate bills were typically repealed in urban areas by the time these state-level abolition policies were passed, so state-level rate bill abolition should not have any impact on urban attendance. Estimated coefficients for urban areas are close to zero and statistically insignificant, in stark contrast to that in rural areas.[21] The null results of the placebo test suggest that state-level rate bill abolition is truly driving the results observed in rural areas.

Column 3 of Table 3 shows the results of my primary specification including only children at least ten years old when they are observed in the Census. The coefficient is qualitatively similar to that estimated on the full primary-school-age population and statistically significant using both critical values from the t-distribution with $G - 1 = 26$ degrees of freedom and the wild bootstrap. This indicates that positive effects on attendance persist for children later in primary school, as opposed to being concentrated in younger children.[22] As was previously mentioned, primary schooling is effectively a prerequisite to secondary school. Therefore, by increasing school attendance for children between the ages of 10 and 14 that were closest to secondary school, this policy provided the foundation for the "high-school revolution" in the early 20th century, which has been documented by Goldin (1998).

et al. (2017) and MacKinnon & Webb (2017).

[21]This supports the identifying assumption that rate bill abolition did not coincide with additional attendance-promoting policies, as is discussed in Section 4.1.

[22]Note that many younger children were already attending schools prior to these policies, as is shown in Figure 2. If one is to think of the marginal student that drops out of school to work, this would be more likely for students at least ten years old than those younger students, especially if the older individuals had already attained some level of education prior to entering the labor force.

Similarly, Column 1 of Table 4 shows the results of my primary specification including only children of lower socio-economic status.[23] These results are qualitatively similar to those presented in Table 3. However, these estimates are statistically different from those for only high socio-economic status individuals, presented in Table A.3.

1.5.2 The American South

The American Civil War (1861-1865) occurred during the time period observed here. While no Decennial Censuses were collected while the conflict was ongoing, it is possible that school attendance in southern states was affected in the aftermath of the American Civil War, either negatively due to the destruction of infrastructure or positively due to changes in norms, opportunities for farm labor, and returns to education. Many southern states passed rate bill abolition laws during the same decade as the conflict, further complicating the issue. On the other hand, the American Civil War should not affect observed attendance in northern states because it was not ongoing during any Census year. While attendance everywhere may have changed during the war, only some states experienced severe destruction and societal change that would have had long-lasting effects on school attendance observed in 1870 or 1880.

In order to avoid the potential confounds of the American Civil War, I replicate the estimation of Equation 1 for the 18 states that did not join the Confederacy and had abolished slavery prior to 1860.[24] Results are presented in Table 5 and are qualitatively

[23]Low socio-economic status is defined as the maximum occupation score in the household under 20. Occupation score is calculated as the median income, in hundreds of 1950 US dollars, for a given occupation in 1950. For reference, a farmer has an occupation score of 14, below the threshold for low socio-economic status as I have defined it. Approximately 60% of rural children are classified as low-SES by this metric.

[24]Excluded states are Arkansas, Florida, Georgia, Louisiana, Maryland, Missouri, South Carolina, Virginia, and West Virginia.

similar to those results presented for the full sample of 27 states.[25] I still find large, statistically significant increases in rural attendance, both for children observed between the ages of 5 and 14 and between the ages of 10 and 14, as a result of the rate bill abolition policies and no such increases in urban areas.[26]

1.5.3 Compulsory Schooling Laws

Some states also passed early forms of compulsory schooling laws in the late 19th century, requiring students to attend school up to a certain age or for a certain number of years. Research from Landes & Solmon (1972) shows that these laws were not effective at increasing attendance, both because these laws were not strictly enforced and were not binding to large swaths of the population that already attended scholl beyond what was required. Due to lack of data on compulsory schooling in some states, I do not control for CSL's in the primary specification. This choice allows me to maintain 27 clusters rather than limiting the sample to 15 states.

I check robustness on the restricted sample to account for their passage using two separate methods. First, I restrict my sample to those 15 states with CSL information and estimate the following equation:

$$Y_{aist} = \alpha + \beta_1 Treatment_{ast} + \beta_2 PartialTreatment_{ast} + \beta_3 CSL_{st} + \gamma_{as} + \delta at + \varepsilon_{aist} \tag{1.2}$$

[25]Estimated coefficients are slightly lower when southern states are excluded from the analysis. I believe that this is driven by either the larger magnitudes of rate bills in the South, the higher prevalence of rate bills in the South, or a combination of those two factors.

[26]The difference between low- and high-SES individuals (presented in Table A.4 and A.5, respectively)

Equation 2 is equivalent to Equation 1 with an additional control for whether or not the state had passed a compulsory schooling law prior to the Census year in which the individuals are observed. Results for this adjusted analysis are displayed in Column 3 of Table 6.

Second, I restrict my sample to the 13 states with known compulsory schooling legislation dates after 1870, [27] truncate my data to only include the 1850 through 1870 Decennial Censuses, and again estimate Equation 1. Results for this adjusted analysis are displayed in Column 4 of Table 6 and are qualitatively similar to those from the primary specification. No matter how compulsory schooling laws are accounted for, the results are qualitatively similar to the estimation of Equation 1 on all four Census years, either on the entire sample of 27 states or the restricted sample with known CSL information.

1.5.4 Other Robustness Checks

I consider an alternative empirical specifications in which partially-treatment individuals are considered by their age when rate bills are abolished rather than as a single category for all children between the ages of 7 and 10. Formally, I estimate the following equation:

is diminished when the South is excluded.

[27] Only Massachusetts and Vermont is known to have had a compulsory schooling law prior to 1870.

$$
\begin{aligned}
Y_{aist} = \alpha + & \beta_1 \mathbb{1}\{\text{year law passed}_s - t + a \leq 6\} \\
& + \beta_2 \mathbb{1}\{\text{year law passed}_s - t + a = 7\} \\
& + \beta_3 \mathbb{1}\{\text{year law passed}_s - t + a = 8\} \\
& + \beta_4 \mathbb{1}\{\text{year law passed}_s - t + a = 9\} \\
& + \beta_5 \mathbb{1}\{\text{year law passed}_s - t + a = 10\} \\
& + \gamma_a s + \delta_a t + \varepsilon_{aist}
\end{aligned}
\tag{1.3}
$$

Results are presented in Table A.6. This specification is slightly more flexible than Equation 1 in that it allows for different treatment effects for those individuals that were already school-aged when these rate bill abolition policies came into effect and are therefore considered partially-treated. However, the disadvantage of this specification is that there are very few individuals in each partially-treated group, since so many individuals are either fully-treated or untreated. Therefore, the coefficients for each age of partial treatment are imprecisely estimated.

It is possible that households may migrate from states that allow tuition charges to those that have banned rate bills. To account for this possible form of selection into treatment, I drop all individuals that do not reside in their state of birth (approximately 15% of all observations) and find qualitatively similar results, which are shown in Table A.7. The pattern of positive, statistically significant effects in rural areas and null effects in urban areas continues for those individuals that reside in their state of birth, indicating that interstate mobility does not affect the results from the primary specification.

1.6 Economic and Historical Significance

The increased primary school attendance here corresponds to an extra 0.72 years of average education for the rural population.[28] This estimate may understate the true effect on educational attainment for three reasons. First, estimated coefficients may be attenuated towards zero for the reasons discussed in Section 4.1. Specifically, the inclusion of older individuals in the control group and the assumption that rural areas charged rate bills until states abolished their use may lead to the inclusion of control observations in the treatment group. Second, the Decennial Census asked whether or not an individual had attended school at any time in the previous year, suggesting that individuals attending part-time would answer yes to this question. Given the structure of the data, I can only observe the changes at the extensive margin; the 7.2 percentage point increase could be shared between individuals switching from no attendance to part-time attendance or from no attendance to full-time attendance. However, these policies could also induce shifts at the intensive margin from part-time attendance to full-time attendance, which would not be picked up in the estimation step because those individuals would have already answered "yes" regardless. As long as the shift from no attendance to part-time attendance is less than the shift from part-time attendance to full-time attendance, the 0.72 additional years is an underestimate. Third, only primary school age children are considered in this analysis. It is possible that finishing primary school would induce some individuals to also receive either a secondary or post-secondary education. Unfortunately, only attendance is observed, and there is no data on education attainment for this time period. Therefore, estimates are based entirely on the flow of education, rather than the observed stock when an individual enters the labor force,

[28] Individuals are 7.2 percentage points more likely to attend for any year in a 10-year period. $0.072 \times 10 = 0.72$

which would account for any cumulative gains of education beyond primary school. This continuation value is an indirect but potentially important aspect of the impact of these policies on educational attainment.

Increased education has large, direct effects on wages in contexts such as this. Duflo (2001) uses school construction in Indonesia in the 1970's to estimate returns to education and finds that an additional year of education increased income by between 6.8 and 10.6 percent. This setting corresponds well to the 19th century United States because the population of Indonesia at the time was primarily rural, largely involved in agriculture, and these policies considered in both cases revolve around increasing the fraction of the population with a primary school education. Based on these results from Duflo (2001), an additional 0.72 years of educational attainment for rural individuals would translate to between an 4.9 and 7.6 percent increase in rural income as a result of rate bill abolition.

In a similar vein, Goldin & Katz (2000) estimate returns to education based on the 1915 State Census of Iowa, which was extremely rare in that it asked not just about contemporary school attendance but also educational attainment. The authors find that returns to education in the agricultural sector were approximately 3.8 percent per year, which translates to a 2.7 percent increase in income based on the additional 0.72 years estimated here. This estimate may be considered a lower bound for the gains from increased education because the marginal year of education at this time was in secondary school. If returns to education are concave, then an additional year would have a higher impact (measured in percent terms) at the primary level than at the secondary level. Additionally, these estimates only account for within-occupation returns; if additional educational attainment also induces workers to switch out of agriculture to more productive occupations, then an annual return of 3.8 percent would understate the true effect of education on productivity.

Education also has indirect effects on the economy and society beyond increased wages. Increased agricultural productivity is important in the history of the development of the broader economy of the United States. Multiple papers (Cao & Birchenall, 2013; Emerick, 2018) have shown that positive shocks to agricultural productivity lead to labor shifting away from agriculture and into other sectors of the economy. Therefore, increased educational attainment that led to higher agricultural productivity would also induce structural change, effectively "freeing up" workers to leave the agricultural sector and instead work in the services sector or the burgeoning industrial sector at the time.

In addition, high-education workers are well-suited to non-agricultural jobs. Although one might think of typical industrial jobs in the 19th century as utilizing "unskilled" labor, workers still had to be educated to properly operate the new machinery and techniques that were being incorporated into the industrial sector. Goldin & Katz (2000) shows that workers with higher education were more likely to hold non-agricultural jobs, indicating that additional education prepared future workers for other sectors as well as agriculture.

Increased education can also lead to higher rates of innovation in an economy. In a stylized setting, Toivanen & Väänänen (2016) find that additional technical universities in Finland led to more patents being filed in the United States by Finnish inventors. Although this is not directly comparable to the 19th century United States, it is reasonable to assume that increased education would lead to higher rates of invention and greater research and development, in addition to simply allowing workers in the agricultural and other sectors to make better use of contemporary modern techniques and technologies. Additional innovation would in turn require greater skills from workers, again highlighting the importance of education and literacy.

Lastly, increased primary school attendance was crucial as the United States became one of the most-educated countries. By 1890, the primary school enrollment rate was approximately 97.1%, higher than every country in the world (Lindert, 2004).[29] These policies also provided the foundation for the "high-school revolution: documented in Goldin (1998). The author herself acknowledges this: "mass secondary schooling in the early twentieth century was made possible because universal elementary education had already spread throughout most sections of the nation" (Goldin & Katz, 2008). Thus, the increased primary school attendance induced by rate bill abolition was a crucial step in the development of the education system in the United States.

1.7 Conclusion

This paper uses United States Census data from the 19th century to investigate the effect of state-level rate bill abolition laws on primary school attendance in rural areas. I find that preventing local school districts from charging tuition and user fees led to increased attendance for children that started school after abolition laws took effect. Cubberley (1919) was indeed correct in asserting that "the [rate bill] charge was small, but it was sufficient to keep many poor children away from the schools." Rate bill abolition was an important policy that increased primary school attendance in the 19th century and helped the United States eventually become the most educated nation in the world during the 20th century.

This analysis is subject to some limitations, primarily driven by the available data on this subject. I rely entirely on state-level policy changes rather than more granular rate bill information at the schools district level. Additionally, I am unable to control for

[29]This figure considers children between the ages of 5 and 14 enrolled in either public or private schools. 85.7% of children between the ages of 5 and 14 were enrolled specifically in public schools, which was

inter- and intra-state mobility, especially that which may be caused by changes in rate bill policies over time. This paper also relies on attendance as it is marked in the 1850 through 1880 Censuses, which is an indicator variable rather than a measure of how much an individual has attended school both in the past year and throughout their childhood. Lastly, I am unable to comment on overall educational attainment other than back-of-the-envelope calculations, since the United States Census did not ask about completed educational attainment until 1940, well after rate bills were abolished throughout the country.

Despite these limitations, this paper contributes to the literature on the history of education in the United States. Previous studies (Denison, 1985; Goldin, 1998) have shown that increased emphasis on education in the early 20th century contributed to the economic success of the United States, and I find that rate bill abolition in the 19th century increased attendance even before such policies as compulsory schooling laws and the standardization of the American high school, initiating the expansion and diffusion of access to education.

This research is relevant to today in two specific contexts. Some countries in the developing world are still undergoing or are yet to adopt these policies implemented in the United States during the 19th and 20th centuries. Places like India have recently abolished tuition payments for public schools, and some areas in Africa continue to charge students to attend public school. These countries have much lower levels of primary school attendance than the rest of the world, in some cases lagging behind the United States by over a century. In the United States, post-secondary education may provide an interesting parallel for primary schools in the 19th century. Post-secondary education is by no means compulsory, just as primary schooling was not compulsory when rate bill

also the highest rate in the world.

abolition was implemented, and community colleges charge low levels of tuition that are dwarfed by marginal returns to education (Marcotte *et al.*, 2005; Marcotte, 2019). The results of this paper suggest that abolition of low tuition requirements and other fees may lead to large increases in primary school attendance in developing countries and post-secondary attendance in the United States.

1.8 Tables

State	Rate Bill Abolition	Compulsory Schooling
New Hampshire	1789	1871
Maine	1820	1875
Massachusetts	1827	1852
Delaware	1829	
Pennsylvania	1834	1895
Louisiana	1847	
Wisconsin	1848	1879
Indiana	1852	1897
Ohio	1853	1877
Illinois	1855	1883
Iowa	1858	1902
West Virginia	1863	
Vermont	1864	1867
Maryland	1865	
Missouri	1866	
New York	1867	1874
California	1867	
Connecticut	1868	1872
Rhode Island	1868	1883
South Carolina	1868	
Arkansas	1868	
Michigan	1869	1871
Florida	1869	
Georgia	1870	
Virginia	1870	
New Jersey	1871	1875
Utah	1890	

Table 1: Rate Bill Abolition Dates. Source: Go (2009)

Table 2: Summary Statistics in 1850

	Before 1850	1850-1859	1860-1869	1870 or later
% attending school	75.2	63.9	67.0	43.6
% female	49.3	48.9	49.2	49.0
% urban	17.7	5.3	14.9	6.5
age	9.3	9.2	9.3	9.3
N	1,059,576	1,066,815	1,367,436	431,562

Source: 1850 Census, Ruggles *et al.* (2019)

The above table gives relevant summary statistics for school-age children in the 27 states considered here in 1850. Columns are organized by the decade in which the state abolished rate bills for students. I compare these groups across some observable dimensions.

Unfortunately, literacy was not recorded in the 1850 Census. Otherwise, it would be included in this table.

Table 3: Staggered Adoption Differences in Differences

	Rural	Urban	Rural 10+	Urban 10+
	prob(att)	prob(att)	prob(att)	prob(att)
Full Treatment	0.0720**	0.0003	0.0790**	0.0147
($age < 7$	(0.0202)	(0.0224)	(0.0232)	(0.0226)
when law passed)	[0.029,0.127]***	[-0.049,0.068]	[0.031,0.094]***	[-0.042,0.086]
Partial Treatment	0.0367	0.0015	0.0425	0.0140
($11 > age > 6$	(0.0209)	(0.0146)	(0.0219)	(0.0136)
when law passed)	[-0.008,0.084]	[-0.023,0.047]	[-0.006,0.094]	[-0.010,0.060]
N	17,610,766	5,704,146	8,398,881	2,636,717

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Standard errors clustered at the state level in parentheses. Critical values are obtained from the t-distribution with 26 degrees of freedom. Confidence intervals obtained from the Wild Bootstrap are shown in brackets. Column 1 shows the results from Equation 1 estimated on rural individuals between the ages of 5 and 14. Column 2 shows the results from Equation 1 estimated on urban individuals between the ages of 5 and 14. Column 3 shows the results from Equation 1 estimated on rural individuals between the ages of 10 and 14. Column 4 shows the results from Equation 1 estimated on urban individuals between the ages of 10 and 14. All regressions include state-by-age fixed effects and Census year-by-age fixed effects, as well as a constant term. No other controls are included.

Table 4: Staggered Adoption Differences in Differences - Low SES Only

	Rural prob(att)	Urban prob(att)	Rural 10+ prob(att)	Urban 10+ prob(att)
Full Treatment	0.0783***	-0.0095	0.0870**	0.0033
($age < 7$	(0.0200)	(0.0208)	(0.0243)	(0.0245)
when law passed)	[0.034,0.129]**	[-0.054,0.054]	[0.034,0.155]***	[-0.053,0.074]
Partial Treatment	0.0386	-0.0003	0.0457	0.0105
($11 > age > 6$	(0.0216)	(0.0193)	(0.0230)	(0.0198)
when law passed)	[-0.010,0.088]	[-0.035,0.063]	[-0.008,0.101]	[-0.025,0.076]
N	10,992,025	754,181	5,294,987	361,536

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Standard errors clustered at the state level in parentheses. Critical values are obtained from the t-distribution with 26 degrees of freedom. Confidence intervals obtained from the Wild Bootstrap are shown in brackets. Column 1 shows the results from Equation 1 estimated on low-SES rural individuals between the ages of 5 and 14. Column 2 shows the results from Equation 1 estimated on urban individuals between the ages of 5 and 14. Column 3 shows the results from Equation 1 estimated on rural individuals between the ages of 10 and 14. Column 4 shows the results from Equation 1 estimated on urban individuals between the ages of 10 and 14. All regressions include state-by-age fixed effects and Census year-by-age fixed effects, as well as a constant term. No other controls are included.

Table 5: Staggered Adoption Differences in Differences - Excluding American South

	Rural prob(att)	Urban prob(att)	Rural 10+ prob(att)	Urban 10+ prob(att)
Full Treatment	0.0401*	-0.0181	0.0400*	0.0006
($age < 7$	(0.0190)	(0.0177)	(0.0179)	(0.0190)
when law passed)	[0.003,0.106]*	[-0.053,0.041]	[0.005,0.099]*	[-0.037,0.068]
Partial Treatment	0.0096	-0.0180	0.0092	-0.0026
($11 > age > 6$	(0.0187)	(0.0088)	(0.0183)	(0.0063)
when law passed)	[-0.027,0.074]	[-0.037,0.007]	[-0.026,0.072]	[-0.016,0.014]
N	13,307,070	5,015,349	6,344,541	2,315,178

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Standard errors clustered at the state level in parentheses. Critical values are obtained from the t-distribution with 17 degrees of freedom. Confidence intervals obtained from the Wild Bootstrap are shown in brackets. Column 1 shows the results from Equation 1 estimated on rural individuals between the ages of 5 and 14, excluding those children in the American South. Column 2 shows the results from Equation 1 estimated on urban individuals between the ages of 5 and 14. Column 3 shows the results from Equation 1 estimated on rural individuals between the ages of 10 and 14. Column 4 shows the results from Equation 1 estimated on urban individuals between the ages of 10 and 14. All regressions include state-by-age fixed effects and Census year-by-age fixed effects, as well as a constant term. No other controls are included.

Table 6: Staggered Adoption Differences in Differences - Compulsory Schooling Laws

	Rural	Urban	CSL on RHS	1850-1870
	prob(att)	prob(att)	prob(att)	prob(att)
Full Treatment	0.0359	-0.0212	0.0515*	0.0511*
($age < 7$	(0.0185)	(0.0173)	(0.0185)	(0.0208)
when law passed)	[0.001,0.100]*	[-0.055,0.037]	[-0.001,0.075]	[0.006,0.119]*
Partial Treatment	0.0047	-0.0194	0.0095	0.0034
($11 > age > 6$	(0.0174)	(0.0091)	(0.0192)	(0.0198)
when law passed)	[-0.031,0.065]	[-0.040,0.006]	[-0.031,0.078]	[-0.049,0.075]
N	12,999,942	4,883,107	12,999,942	8,720,046

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Standard errors clustered at the state level in parentheses. Critical values are obtained from the t-distribution with 14 degrees of freedom (12 degrees of freedom for Column 4). Confidence intervals obtained from the Wild Bootstrap are shown in brackets. Column 1 shows the results from Equation 1 estimated on rural individuals between the ages of 5 and 14 only in the 15 states for whom data on compulsory schooling laws is available. Column 2 shows the results from Equation 1 estimated on urban individuals between the ages of 5 and 14. Column 3 shows the results from Equation 2 estimated on rural individuals between the ages of 5 and 14. Column 4 shows the results from Equation 1 estimated on rural individuals between the ages of 5 and 14 using only data from 1850 through 1870, excluding Massachusetts and Vermont. All regressions include state-by-age fixed effects and Census year-by-age fixed effects, as well as a constant term. No other controls are included except the indicator for a state having passed a compulsory schooling law in Equation 2.

1.9 Figures

1.9.1 Map of Included States

Figure 1

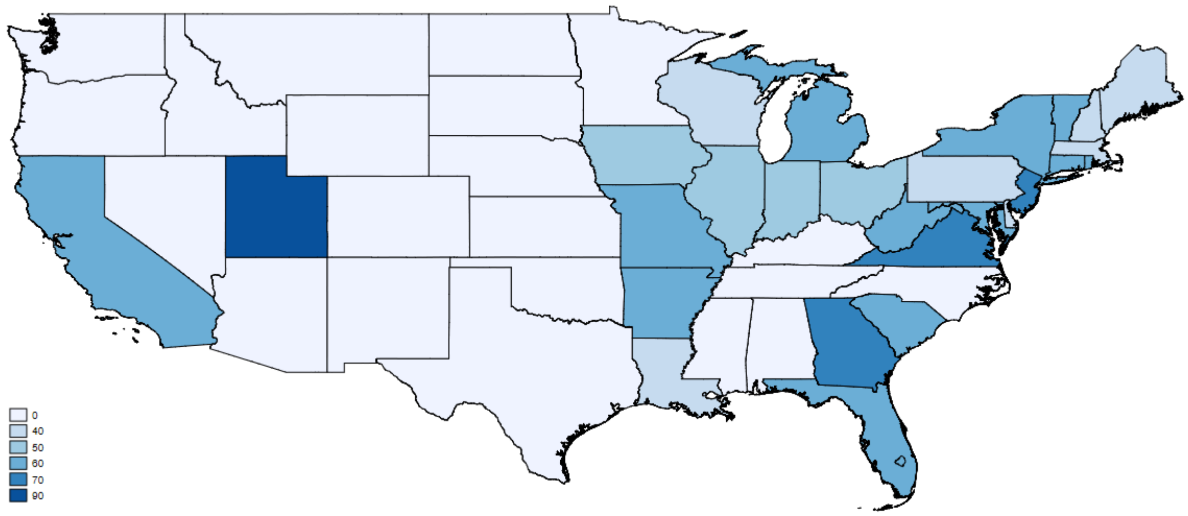

Source: Go (2009)

1.9.2 Attendance Rates by Age

Figure 2

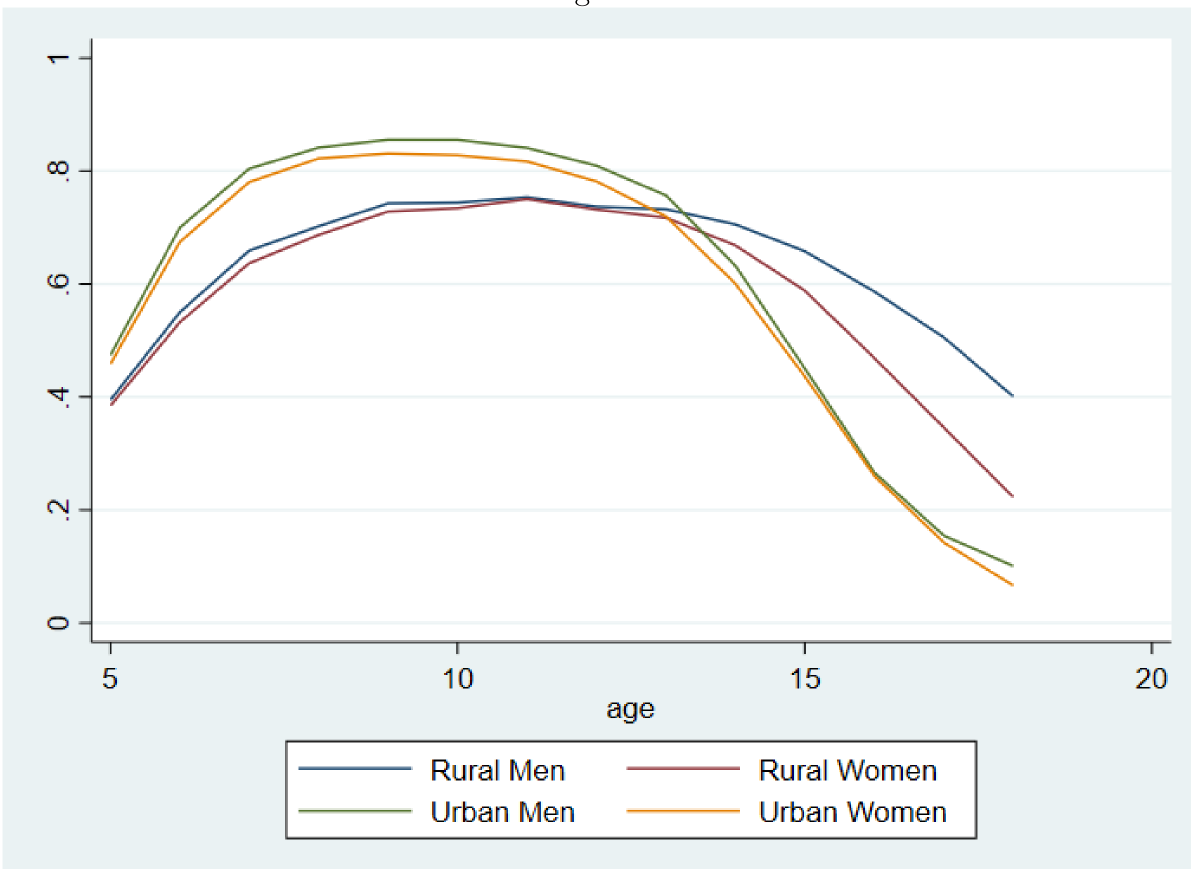

Attendance Rates by Age, broken into subgroups. Source: 1850 Census

Chapter 2

Urbanicity and Fertility: Evidence from Refugees to Canada

2.1 Introduction

It is well-established that urban areas have lower levels of fertility than rural areas. However, it remains unclear to what extent this relationship is direct (driven by factors related to the urban environment) as opposed to indirect (driven by differences in who chooses to live where). This dichotomy is well summarized by the following quote: "two competing hypotheses are discussed in the literature: the compositional and the contextual. The compositional hypothesis suggests that fertility levels vary between places because different types of people live in different settlements, whereas the contextual hypothesis suggests that factors related to immediate living environment are of critical importance" (Kulu & Washbrook, 2014). This paper attempts to decompose the direct and indirect effects of urbanicity by comparing fertility rates among refugees who are exogenously placed throughout Canada, finding that rural-urban fertility differences persist among those refugees and are therefore primarily driven by contextual factors rather than compositional factors.

Rural-urban differences in fertility have been documented since at least United Nations (1953). Kuznets (1974) summarizes the data through the mid-1960's, documents differences across countries , and notes that migration from rural to urban areas is not correlated with changes in birth rates, suggesting that there is no causal effect of urbanicity on fertility. Lerch (2018) also examines rural-urban fertility differences in the developing world over time, emphasizing the urban fertility decline and the role of rural-to-urban migration in mitigating fertility differences across type of residence. The author also describes the potential importance of where young women are "socialized" in determining preferences for fertility, one potential cause of the rural-urban fertility differences. However, rural-to-urban migration can only illuminate aggregate trends in fertility rather than fertility differences, which are instead driven by urban fertility declines that outpace those in rural areas. Lerch (2018) focuses on developing nations, whereas developed countries have long-standing and persistent rural-urban fertility differences that imply they are not functions of a "transitional period" but instead an equilibrium object.

Kulu (2013) utilizes longitudinal data from Finland and argues that contextual factors play some role in rural-urban fertility differences. Of the potential mechanisms, the author specifies three specific possibilities: higher prices in urban areas, higher opportunity costs of working in urban areas, and differences in cultural norms between rural and urban areas. However, the model fails to control for unobservable factors that may drive composition effects. Specifically, preferences for fertility may drive the location choice for households in that they may sort into areas orthogonal to any observable characteristics such as education and race. These results and limitations are also true of

Kulu & Washbrook (2014), which extends the previous results to the United Kingdom using vital statistics data. Simon & Tamura (2009) and Clark (2012) both utilize panel data and find that higher housing prices are correlated with lower fertility. However, both papers rely on endogenous prices and household choices rather than exogenous variation, and the results may therefore be driven by selection (either observable or unobservable) rather than the direct effect of urbanicity.

There is also an extensive literature on the causal relationship between income and fertility, dating back to Becker (1960). Black *et al.* (2013) discusses the role of urbanicity in determining both wages and prices, with the eventual goal of shedding light on the quality-quantity trade-off. They extend Roback (1982) to highlight the importance of accounting for urbanicity when exploring the income-fertility relationship. Finally, the authors evaluate the impact of income on fertility using an exogenous shock in coal prices during the 1970's in Appalachia and find that children are indeed normal goods.

Using data from the 2016 Census Long Form Questionnaire in Canada, I compare fertility rates among refugees that live in rural and urban areas, for whom the compositional effect should be mitigated. I provide similar information for both non-refugee immigrants and native-born Canadians, for whom both compositional and contextual factors could lead to rural-urban fertility differences. In order to account for both selection on observable characteristics and the potential for non-random placement based on those same observables, I include fixed effects for age, education level, number of children older than one year old, and origin (either province or foreign country of birth). In order to account for differences in socialization and marital patterns across urbanicities, I focus on women who were married prior to their arrival in Canada.

I exploit the exogenous placement of refugees throughout Canada to isolate the causal, or contextual, effects of urbanicity on fertility. When refugees are admitted into Canada, they are distributed across the country by the Resettlement Operations Center in Ottawa (ROC-O). ROC-O takes into account various refugee characteristics, such as language and medical needs, and potential destination characteristics, such as existing ethnic communities and the availability of resettlement services. Crucially, the decision is made entirely by the resettlement office and not by refugees, so the final placement of refugees in Canada is exogenous to the refugee families.

Canada has a long history of accepting refugees.[1] Approximately 1 million individuals, representing 2.4% of Canada's population and over 10% of all immigrants living in Canada, entered the country as refugees.[2] Canada accepts the second-most refugees in the Americas, approximately one third as many refugees as the United States despite a population only one tenth as large.[3]

Previous research has used refugees for quasi-experimental evidence to answer questions in other fields of economics. Two especially influential strands are Card (1990) on the Mariel Boatlift and Foged & Peri (2016) on the long-run impacts of immigration on wages in Denmark. Both of these papers primarily study the effects of refugees on existing communities. In contrast, this paper exploits the exogenous placement of refugees across urbanicities to investigate the potential effects of that urbanicity on the fertility decisions of the refugees themselves.

Refugees are individuals forced to flee a home country and permanently settle elsewhere. The United Nations Refugee Agency identifies individuals for resettlement,

[1] For an anecdotal description of Canada's history of refugees, visit https://www.canada.ca/en/immigration-refugees-citizenship/services/refugees/canada-role/timeline.html

[2] Based on calculations from the Public Use Microdata Files from the 2016 Canadian Census.

[3] According to data from the United Nations High Commissioner for Refugees, available from the World Bank at

at which point either the Canadian government or private groups sponsor those refugees to travel to Canada and receive basic services. In addition to granting residence and protection in Canada, refugees benefit from various resettlement services including up to one year of income support, help finding temporary or permanent housing, clothing, food, and information about Canada in general and local services specifically. In some cases, sponsors may provide assistance in one form or another for up to three years.

I find that rural-urban fertility differences persist among those refugees, both with and without fixed effects included, suggesting that these differences are primarily driven by contextual factors rather than compositional factors. The inclusion of fixed effects decrease the rural-urban difference by approximately 27%. This differences remains similar to that for non-refugee immigrants.

The rest of the paper proceeds as follows: Section 2 introduces the data utilized in this paper. Section 3 describes the empirical strategy. Section 4 highlights the key findings. Section 5 concludes.

2.2 Data

I utilize the 2016 Canadian Census 25% sample of the Long Form Questionnaire. The 2016 Canadian Census was the first of its kind in Canada, as well as one of the first Censuses overall, to ask about specific refugee status as a potential admissions category for immigrants into Canada. These data were accessed at the Research Data Centre in Vancouver, British Columbia. The dataset includes information on various facets of life, including family, demographics, geographic mobility, and education for almost 10

https://data.worldbank.org/indicator/SM.POP.REFG million Canadians.I define individuals as native Canadians, refugees, and non-refugee immigrants (NRI's) to partition the set of possible origins and admissions categories. I restrict the sample to women between the ages of 15 and 45 (inclusive) and focus on urban areas primarily as the three largest cities in Canada: Toronto, Montreal, and Vancouver. Summary statistics from this section are from the Individual-Level Public Use Microdata Files (PUMF) rather than the 25% Sample, as the release of descriptive statistics is restricted.

About 9.1% of households in Canada have a child aged either zero or one year old. This fraction is slightly larger (9.3%)for native-born Canadians, which constitute the largest of the three groups. Implied fertility rates are lower for refugees (8.1%) and non-refugee immigrants (7.9%). This could reflect the fact that both immigrants and refugees may have recently arrived in a new country and be adjusting to a new area and culture, which could affect their fertility decisions.

Approximately 36% of the Canadian population lives in one of the three largest cities, but this figure varies greatly across the three categories as I have defined them. Approximately 30% of native Canadians live in the three largest cities, compared to about 60% of refugees and about 65% of NRI's. Non-refugee immigrants are also more likely than refugees to live in truly rural areas, but refugees are far more likely to live in cities other than the top three (for example: Calgary, Edmonton, Ottawa, etc.).

Geographic mobility of immigrants and refugees are low and similar to each other. In both cases, 75% of individuals that arrived in 2010 live in the same CSD (Census Subdivision) as 5 years ago. That number rises to over 90% for one year mobility for 2014 arrivees.[4] This may suggest limited geographic mobility in that it is difficult

[4]These figures excludes immigrants and refugees that arrived fewer than six and fewer than two years

for immigrants and refugees to move. However, it may also indicate that non-refugee immigrants (refugees) initially move to (are exogenously placed in) favorable locations in which they prefer to stay rather than relocate. These figures are similar for individuals that arrived in Canada without the timing restrictions, suggesting that there is not a significant initial migration upon arrival to address any misallocation in initial destinations. In either case, the initial conditions are sticky, and there is very little movement within Canada once refugees and immigrants migrate.

Non-refugee immigrants to Canada are very positively selected with respect to education; approximately 45% of this population has at least a bachelor's equivalent degree. This starkly contrasts the other two groups, refugees and native Canadians, which are very similar in terms of education, for whom the fraction with at least a bachelor's degree is 25% and 36%, respectively.

2.3 Empirical Strategy

I estimate the following equation on the full sample to establish the observed rural-urban fertility differences:

$$P(child \quad born)_{io} = \alpha + \beta Urban_i + \varepsilon_{io} \tag{2.1}$$

where $P(child \quad born)_{io}$ is an indicator equal to 1 if a woman has at least one child born in the past year,[5] $Urban_i$ is an indicator equal to 1 if the individual lives in one of the three

ago, respectively.

[5]This indicator outcome variable is chosen to abstract from twins, triplets, etc. which are not "chosen" in the same way.

largest metropolitan areas of Canada, and ε_{io} is an error term clustered at the "origin" level to account for potential cultural differences that might affect fertility decisions even after immigration to Canada.[6]

Second, I include various fixed effects to control for many observable ways in which urban potential mothers may differ from their rural counterparts. I augment Equation 1 to the following:

$$P(child \quad born)_{aeiko} = \alpha + \beta Urban_i + \gamma_a + \delta_e + \lambda_k + \eta_o + \varepsilon_{aeiko} \tag{2.2}$$

where γ_a is an age fixed effect, δ_e is an education fixed effect, λ_k is a number-of-children fixed effect,[7] and η_o is an origin fixed effect. The coefficient of interest β represents the difference between rural and urban fertility rates among observably-similar women. For refugees that are exogenously placed across Canada, this difference highlights the potential causal effect of urbanicity on fertility.

2.4 Results

Column 1 of Table 1 shows the pure differences in fertility rates between rural and urban women for refugees. Columns 2 and 3 show the same for non-refugee immigrants and native-born Canadians, respectively. The coefficient for native-born Canadians is especially noteworthy, indicating that urban native-born Canadians actually have more children than their rural counterparts when no controls are included. This is surprising,

[6]The origin here refers to the place of birth, either a province or territory within Canada or a foreign country.

[7]The number of children used here is equal to the number of children living in the census family over the age of 1.

runs counter to observed rural-urban fertility differences in general, and may be worthy of further exploration in the future.

Column 4 of Table 1 shows the rural-urban differences in fertility among refugees after accounting for various observable characteristics. Columns 5 and 6 show the same for non-refugee immigrants and native-born Canadians, respectively. The coefficient on an indicator for living in an urban area drops by 27% for refugees. This decrease suggests that refugees allocated to rural and urban areas differ on various observable characteristics, in which case the inclusion of these variables on the right hand side is crucial for correct comparison across these two groups.

The coefficient on non-refugee immigrants remained largely unchanged, and the coefficient on native-born Canadians flips entirely and becomes statistically significant in the opposite direction, highlighting the importance of controlling for these various characteristics in fertility analysis.[8]

Based on Column 4 of Table 1, urbanicity causes a 1.1 percentage point drop in the probability of having a child in the past year. This result is statistically significant at the 1% level and reflects a 9% decrease from that probability among refugees in rural areas. The fertility rate among refugees in rural areas is higher than that of non-refugee immigrants and native Canadians in rural areas, so this causal impact reflects a larger fraction of their fertility than it does for refugees.

[8]Accounting for geographic mobility has no effect on the estimated coefficients. Columns 1 through 3 of Table A.2 replicate Columns 4 through 6 of Table 2, and Columns 4 through 6 of Table A.2 show the results when the sample is restricted to only women that live in the same Census Metropolitan Area as five years ago. Results are exactly the same on the full and restricted samples. Although this does not

2.5 Conclusion

This analysis suggests that causal factors are primarily responsible for rural-urban differences in fertility. These rural-urban differences persist for refugees, for whom placement within Canada is exogenous, even after accounting for various observable characteristics. Compositional factors, either based on observable or unobservable characteristics, should not play any role in the estimated coefficient of urbanicity on fertility, since refugees do not choose where in Canada to migrate to. Therefore, the observed differences reflect direct causal impacts of urbanicity on fertility and suggest that these causal factors drive observed differences.

Although the primary focus of this paper is on refugees, the results raise interesting questions for non-refugee immigrants and native Canadians. The estimated coefficients for non-refugee immigrants are lower than those for refugees, even though this group should be subject to both causal and compositional effects of urbanicity. This may be driven by inherent differences between refugees and non-refugee immigrants, specifically that non-refugee immigrants are positively selected (relative both to refugees and native Canadians) and may therefore be less affected by the direct effects of urbanicity.[9]

There are various limitations with this analysis. First, the estimated coefficients are unable to account for any differences in socialization among migrants to Canada. The existing literature has suggested that the social norms an individual experiences as they grow up may affect their fertility preferences. However, although fixed effects

fully account for the possibility that refugees or non-refugee immigrants move after arriving in Canada, the fact that there is little geographic mobility immediately upon arrival and recent the exclusion of recent movers does not affect coefficients is reassuring.

[9]Specifically, immigrants may be better able to afford higher costs of living. Similarly, if the primary earner in an immigrant household earns more than that of a refugee household, than a potential secondary earner may choose not to work. Therefore, the higher opportunity cost of working for potential mothers may affect non-refugee immigrants less than refugees.

for country of birth are included on the right-hand-side of Equation 2, this variable does not fully account for difference in socialization. Further research may exploit the exogenous placement of refugee children to Canada, rather than married refugees, to further investigate this potential channel through which contextual factors may affect rural-urban fertility differences.

In addition, the refugee placement does not provide as clean variation as in Foged & Peri (2016). Although the placement of refugees within Canada is exogenous, private sponsors and the government may still take into account various observable characteristics in their decision-making process. For example, refugee households with children may be more likely to be placed in rural areas where the cost of living is lower. Therefore, the groups that eventually move to rural and urban areas are not necessarily ex-ante identical, indicating a departure from the ideal natural experiment. However, the lack of geographic mobility once in Canada does suggest that initial destinations really are sticky and that refugees do not internally migrate for fertility (or other) purposes.

Lastly, this analysis focuses on refugees moving to a new country fleeing persecution. They are by construction adjusting to a new culture. While this is helpful for the purposes of their exogenous placement within Canada, their situation does not necessarily reflect that of most people. Further analysis might exploit quasi-experimental variation in urbanicity in a different context and add to the literature on this topic.

Table 1: Rural-Urban Fertility Differences

	(1)	(2)	(3)	(4)	(5)	(6)
	Refugee	NRI	Native	Refugee	NRI	Native
Urban	-0.015	-0.008**	0.012**	-0.011**	-0.009***	-0.010***
	(0.009)	(0.003)	(0.003)	(0.004)	(0.002)	(0.001)
Rural Mean	0.122	0.104	0.098	0.122	0.104	0.098
% of Rural Mean	12.3%	7.7%	-12.2%	9.0%	8.7%	10.2%
Age Fixed Effects				X	X	X
Educ Fixed Effects				X	X	X
# Kids Fixed Effects				X	X	X
POB Fixed Effects				X	X	X
Clustered?	POB	POB	POB	POB	POB	POB

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Standard errors clustered at the "Place of Birth" level in parentheses, where the "Place of birth" refers to either the province or foreign country or birth. Column 1 shows the rural-urban fertility differences among refugees without accounting for any observable characteristics. Columns 2 and 3 show the same differences for non-refugee immigrants and native Canadians, respectively. Column 4 shows the rural-urban fertility differences among refugees after the inclusion of various fixed effects to control for some observable characteristics. Columns 5 and 6 show the same differences for non-refugee immigrants and native Canadians, respectively. Estimated coefficients drop by approximately 27% for refugees when fixed effects are included, are similar across specifications for non-refugee immigrants, and flip signs for native-born Canadians.

Chapter 3

Do Fans Impact Sports Outcomes? A COVID-19 Natural Experiment

3.1 Introduction

Home teams generally outperform away teams in a wide variety of sports. This stylized fact is well-established and known as home field advantage. Over the course of the 2008-2009 through 2018-2019 season in four top European leagues, home teams outperformed away teams in goal difference by 0.38 goals and win percentage by 17 percentage points.[1] This advantage is driven by various factors that are typically broken down into three main channels: travel fatigue, venue familiarity,[2] and crowd support. It

[1]Home teams win approximately 47% of the time, draw 25% of the time, and lose 29% of the time

[2]In addition, there may be some advantage regarding field preparation and style of play. In particular, FC Barcelona was known in the mid 2010's to keep their grass especially short and somewhat wet, such that short passes were easier to complete. Teams hosting FC Barcelona would sometimes prepare their grass to be long and dry in order to hamper the visitor's short passing style. https://www.dailymail.co.uk/sport/football/article-2600454/Barcelona-complain-UEFA-length-grass-Vicente-Calderon-ahead-Champions-League-showdown-Atletico.html. However, evidence for this

is an empirical challenge to separately identify each effect. While fan attendance is indeed measurable, there are endogeneity concerns regarding the impact of crowd attendance on team performance; better teams typically have bigger stadia and higher attendances, raising concerns about reverse causality.

In this paper, we exploit exogenous variation, driven by the COVID-19 pandemic, in the level of fan attendance. We leverage the fact that four of the top European soccer leagues implemented no-fans policies for the last quarter or so of the 2019-2020 season. These policies cause an exogenous downward shock in attendance levels at soccer games. We find that these no-fans policies decrease the home field advantage in European soccer as measured by goal difference, expected goal difference, and the probability of a home win. These results are robust to inclusion of various fixed effects as well as controls for weather, team strength, and local exposure to COVID-19.

We explore both the causes and implications of this change in home field advantage. Using the expected goals metric, we show that the decrease in home field advantage is driven by home teams playing worse relative to away teams after the implementation of the no-fans policy, as opposed to scoring relatively fewer goals on the same quality and quantity of goal-scoring opportunities. Our results show that the change in home field advantage decreases the probability a home team wins the game by 5.4 percentage points. We find that there is no effect on the probability of a draw between the home team and the away team as decreases in wins led to a direct increase in losses for the home team. This finding suggests that fans have a symmetrically pivotal impact on the outcome of games — the same number of games are shifted from home win to draw as

type of advantage remains anecdotal to this point and is unidentified relative to the venue familiarity effect. Therefore, if one prefers to think of this more generally, then the venue familiarity effect can also include such factors as pitch preparation, the colors in the locker room, and the exact jerseys worn by each team.

are shifted from draw to home loss.

Of all the potential sports leagues with which to examine the impact of fan attendance on home field advantage, the European soccer leagues discussed here provide the best natural experiment. The rules of the game were not materially affected in any way that might plausibly impact home field advantage.[3] While games were delayed between two and three months, no games were cancelled or relocated. This is in stark contrast to both other soccer leagues (Major League Soccer in the United States, European club tournaments such as the UEFA Champions League and the UEFA Europa League), international soccer matches, and North American leagues in other sports (National Basketball Association, Major League Baseball, National Football League, etc.).

This paper contributes to a rich literature focused on identifying the effects of crowd support, travel fatigue, and venue familiarity on the home field advantage. Broadly speaking, the literature on crowd support prior to COVID-19 can be split into two strands. The first strand uses variation in turnout by country and league and generally finds larger advantages for the home team associated with higher fan attendance (Pollard & Pollard, 2005; Clarke & Norman, 1995; Agnew & Carron, 1994). In the second strand, two recent papers leverage quasi-experimental variation in fan attendance to overcome endogeneity concerns. Belchior (2020) uses randomized game times in Brazilian football and finds that fan attendance levels have no effect on home field advantage. Ponzo & Scoppa (2018), on the other hand, provide evidence for the importance of crowd support using same-stadium derbies.[4] In this setting, both teams are familiar with the stadium and do not suffer travel fatigue, but tickets are allocated differently to the home and away

[3]See Table 3.1 for more details on how European soccer leagues responded to the COVID-19 pandemic.

[4]A same-stadium derby is a game between two teams that share a home stadium. The most famous example is AC Milan and Inter Milan, which both share the San Siro stadium. Such games occur twice per year, with each team technically "home" for one of the two games.

teams. Using 20 years of Serie A data and 5 same-stadium derbies, they find that the home team is 15 percentage points more likely to win than the away team, significantly lower than the 25 percentage point advantage in normal matches. We generalize the results of Ponzo & Scoppa (2018) and also find that fans drive a large component of home field advantage in soccer beyond same-stadium derbies in Italy.

Various papers exploit the variation in attendance from COVID-19 to identify the effect of fans on home field advantage. The vast majority of these papers find that the no-fans policies decreased home field advantage (Ferraresi *et al.*, 2020; Fischer & Haucap, 2021; Scoppa, 2021), but some find statistically insignificant effects (Reade *et al.*, 2021; Bryson *et al.*, 2021).[5] Some differences in results can be attributed to heterogeneous effects by country (Benz & Lopez, 2021) and by division (Fischer & Haucap, 2021). We focus on the top divisions in England, Germany, Italy, and Spain, which are the top-ranked leagues in European football, and include a rich set of controls to best isolate the effect of fans on home field advantage in the most important leagues.

This paper differs from the recent literature on home field advantage during COVID-19 in a few important dimensions. First, we examine the effect of the no-fans policy not just on realized goals but also on expected goals. Expected goals are a measure of team performance that accounts for the quantity and quality of shots, which allows us to better isolate the impact of fans on performance. Second, we include controls for various measures of COVID-19 prevalence. Although it is difficult to know ex-ante how COVID-19 would affect home field advantage directly, we control for potential heterogeneous regional effects using data on COVID-19 cases from Carleton *et al.* (2021). Third, we control for differences in weather relative to when these games would have otherwise

[5]Many of these papers focus on changes in referee behavior as the potential mechanism (Endrich & Gesche, 2020; Scoppa, 2021; Reade *et al.*, 2021; Bryson *et al.*, 2021; Fischer & Haucap, 2021) with most finding that referees are less favorable to home teams without fans.

been played; the hiatus in the 2019-2020 season meant that games originally scheduled for March through May were instead played in May through August. Therefore, the temperature during these games was hotter than would have otherwise been expected, and this may lead to differences in home field advantage. By controlling for geospatial variation in cases and intertemporal changes in temperature, we are able to better account for the direct effect of the pandemic and the delay in the season.

Our results indicate that between 35% and 45% of the home field advantage can be attributed to non-fan factors of travel fatigue and venue familiarity. Oberhofer *et al.* (2010) analyze the effect of distance on home field advantage in the Bundesliga and find that the performance of the away team decreases with distance traveled. Similar effects have been found in the English Premier League (Clarke & Norman, 1995), U.S. National Football League (Nichols, 2014), and Australian Football League (Goumas, 2014), some of which also implicate crossing time zones as one channel through which travel may affect home field advantage. Familiarity also likely plays a role in the residual home field advantage as players on the home team are more comfortable in their own stadium. This has an especially large effect when the field is unusually large or small or has an artificial surface (Dowie, 1982; Barnett & Hilditch, 1993; Clarke & Norman, 1995). To causally identify the effect of familiarity, researchers have used team stadium moves that reduce familiarity for the home team. Evidence on this effect is mixed; Pollard (2002) finds a drop in home field advantage after moving, while Loughead *et al.* (2003) finds no effect. In this paper, we are not able to separately disentangle these two non-fan factors, but our work does imply that crowd support is the largest driver of home field advantage.

This paper is organized as follows: Section 2 provides background information and describes the data used. Section 3 discusses threats to identification. Section 4 explains the empirical strategy. Section 5 highlights the key results of our analysis. Section 6 concludes.

3.2 Background and Data

In this section, we describe the five major European soccer leagues, their response to COVID-19 and the data utilized in the empirical analysis.

3.2.1 European Soccer and COVID-19

The top five ranked professional football leagues in the Union of European Football Associations (UEFA) are the Spanish La Liga, English Premier League, Italian Serie A, German Bundesliga and French Ligue 1. These leagues have dominated both European soccer and global sports revenues in recent history.[6] In 2017-2018, these five leagues generated approximately 17.4 billion dollars (USD) in revenue, which is more than the total GDP of Jamaica and double the total revenue of the National Football League (NFL), highlighting the importance of these leagues relative to other sports and the global economy.

Seasons for all of the major European soccer leagues begin in August and end the following May. In the 2019-2020 season, however, the normal schedules were interrupted when COVID-19 was declared a pandemic by the World Health Organization on March 11, 2020. Amid concerns for player and staff safety, all major European leagues halted play from that date until mid-May at the earliest, with some slight variation by country.[7]

[6] Nearly every winner and runner-up in the most prestigious European club competition, the Champions League, has come from these five leagues from the 1996-1997 through 2020-2021 seasons. The only exception is FC Porto, who won the Champions League in 2003-2004.

[7] Each of these five major leagues paused their seasons within a few days of each other. The German

Four of the five leagues resumed play with a restrictive set of protocols.[8] Many of the policies instituted across the leagues were similar during the season,[9] but with some variation in the lead-up.[10] In addition to testing and quarantine policies, each league allowed for two extra substitutions per match and instituted water breaks in order to alleviate the burdens of the hiatus and warmer weather. No games were cancelled or moved, but all four of the leagues decided to play without any fans in attendance.

A brief summary of the timeline, on-field policy changes, and primary method of travel for each league is shown in Table 3.1. All of the leagues stopped at about the same time, but play resumed at different points of time for each league. The Bundesliga resumed play slightly over 2 months after the initial pause, while the other leagues waited approximately 3 months before resuming play. The mitigation policies were successful, no matches were cancelled or rescheduled, and very few players missed games.[11]

3.2.2 Schedule Construction in European Soccer

Teams in the five soccer leagues discussed in this paper all play what is typically called a "balanced" schedule. Every team plays each other team in the same league exactly twice, once at home and once away. Therefore, each team plays $2 \times (n - 1)$

Bundesliga was both the last league to stop play and the first to resume play. The last game before the hiatus was played on March 11th, and the first game after the restart was played on May 17th.

[8] The exception here is the French Ligue 1, which decided to cancel the rest of the 2019-2020 season rather than resume play after the hiatus.

[9] The similarities across leagues are not surprising since the heads of all five of the major European soccer leagues had weekly meetings starting in April.

[10] Spain, for example, instituted a four phase program in the lead-up to the season, which was slightly different than what was instituted in other leagues.

[11] The most publicized case of a player missing a game was Claudio Pizarro. He missed the game after the restart because his daughter tested positive. It is worth noting that there was an outbreak in a second division team in Germany, but no such flare ups happened in any of the top division leagues considered here.

games in a season, where n is the number of teams in that league.[12] Games against the same opponent are spaced throughout the season. In each of these leagues, a certain team will play each potential opponent once in the first half of the season and once in the second half. These two games are often called "reverse fixtures" since they contain the exact same teams but are played at opposite locations and times in the season. In some leagues, such as the Bundesliga, the order of opponents is exactly the same in each half of the season.[13] For example, if Bayern Munich plays at Borussia Dortmund in each team's first game of the season, then Borussia Dortmund will play at Bayern Munich in the eighteenth game of the season.

Games are randomly allocated subject to some restrictions. These restrictions are in place to prevent long strings of home or away games for any given team and to avoid concurrent home games between two teams in the same metropolitan area. If games were perfectly randomly allocated, then there would occasionally be long strings of home or away games for a given team. If this were the case, it could present problems for analysis; results could be driven by randomness in the schedule rather than any changes in home field advantage. In the last ten matches of each league discussed here, every single team played between four and six home games, which mitigates the concern that the home field advantage in the back part of the season is driven by a few teams with many home games.

[12]There are 20 teams in the English Premier League, Italian Serie A, and Spanish La Liga but only 18 teams in the German Bundesliga and French Ligue 1.

[13]Due to scheduling conflicts arising from other competitions this does not always hold.

3.2.3 Expected Goals

Due to randomness, human error, and occasional moments of athletic brilliance, the realized score of a match is a noisy signal for which team actually played better over the course of 90 minutes. In order to mitigate this noise, we focus on expected goals, or xG, which measure the quantity and quality of each team's chances to score; they have been shown to better predict future performance and more closely track team actual performance than realized goals (Rathke, 2017). Expected goals are calculated by summing the ex ante probabilities that each shot, based on its specific characteristics and historical data, is converted into a goal.[14] For example, if a team has four shots in a game, each with a scoring probability of 0.25, then their expected goals for the match would sum to 1. However, their realized goals could take any integer value from 0 to 4.[15]

Note the important difference between expected goals and realized goals in their respective data generating processes. Realized goals effectively transform the expected goals probabilities, in which p can take on any value between zero and one, into a Bernoulli variable with probability of success p equal to the expected goals metric. This process of "converting" chances into goals is subject to increased randomness, and we therefore focus on expected goals as the more precise measure of team performance.

To determine if fans impact the quantity and quality of chances created, we abstract from realized goals and instead use expected goals for each team as our primary measure of performance. We restrict the panel to years for which the expected goals data are available and estimate the effect of the no-fans policy on home field advantage as measured

[14]For more detailed description of how exactly expected goals are calculated, please refer to https://fbref.com/en/statsbomb/. Chances are considered similar based on distance, angle, goalkeeper and defender position, and shot type (shot with the foot or the head).

[15]The realized goals could in fact be greater than four if the other team scores one or more own goals, which introduces even more noise.

by expected goal difference. Importantly, this allows us to identify the effect of fans on quality of play, as measured through shot quality and quantity, rather than the noisier signal of realized goals.

3.2.4 Data Description

We use match data from the soccer statistics website FBref,[16] which provides information on all matches in the five major European soccer leagues going back at least 10 years. Our panel starts with the 2009-2010 season and thus includes 10 years of data prior to the 2019-2020 season and 15,906 matches overall. The data are slightly more limited for expected goals, which are only available on FBref starting in the 2017-2018 season.

We control for daily weather conditions at the location of the match using European Climate Assessment and Dataset (ECA&D) (Klein Tank *et al.* 2002).[17] Game temperatures are then assigned weather data by matching the closest station-level daily average temperatures.[18]

To measure the direct effect of COVID-19 across space, we use one of the most spatially-resolved global dataset of daily confirmed COVID-19 cases from Carleton *et al.* (2021).[19] Combining the geographic location of each stadium with the geographic regions

[16] FBref.com launched in June 2018 with league coverage for six nations: England, France, Spain, Italy, Germany, and the United States.

[17] Specifically, we use the blended daily mean temperature dataset, which can be found here: https://www.ecad.eu/dailydata/predefinedseries.php.

[18] In Figure A.3.1, we show the distribution of distances between the stadium and station. Nearly 90% of observations have a station within 100 kilometers. Serie A matches tend to be farther from the weather stations, but the conclusions from our analysis are robust to the exclusion of Serie A.

[19] Spatial granularity varies for each country in this analysis. Broken down by the Nomenclature of Territorial Units for Statistics (NUTS): Germany is level 1 (16 states), Spain is level 2 (19 autonomous communities), and Italy is level 3 (107 provinces). For England, the data is at the national level and

in the COVID-19 dataset, we calculate three different measures for the effect of COVID-19 on an area as of March 31, 2020: cumulative cases, new cases, and total cases per capita. For each of these measures, we calculate the difference between the home and away team and will use this as a measure of differences in the effect of the virus in the area of the home team compared to the away team.

3.2.5 Summary Statistics

Table 3.2 presents key summary statistics aggregated across all four of the top five European soccer leagues that resumed play in May and June of 2020. Almost all match data is from the 2009-2010 season through the 2019-2020 season. The lone exception is the expected goals metric, which is only available starting in the 2017-2018 season. As a measure of distance traveled, we calculate the linear distance between the home and away team stadiums.

There are two things worth highlighting from Table 3.2. First, home teams generally score more than away teams, which is reflected in a home minus away goal difference of 0.38. Second, there is a similar advantage to home teams in the home minus away expected goals difference. This highlights that the difference in actual goals reflects differences in goal-scoring opportunities. In other words, this table shows the core of what the home field advantage is — a difference in play between home and away teams, which then manifests in goals and match outcomes.

The average home field advantage, attendance, and distance between teams is different across each of the four leagues. In Table 3.3, we break down the summary statistics by

so we are not able to estimate differential cases by team. Our results are robust to the exclusion of the English Premier League

league. The Bundesliga and Premier League have the highest average attendances, and La Liga and Serie A have higher average distances. The average home field advantage, measured by goal difference, is between 0.34 and 0.46 goals per game. Average home field advantages in expected goals are slightly lower, between 0.25 and 0.37 goals per game.

3.3 Identification

Following the hiatus, no-fans policies were implemented in four major European soccer leagues. However, it is possible that non-fan factors could have impacted home field advantage. In this section, we discuss the potential threats to identification and how we address them.

The most direct and obvious potential confounding factor is that the COVID-19 pandemic may have affected home and away teams differently through infection. This is unlikely due to the low number of positive tests among soccer players. The players available to participate for the last quarter of the 2019-2020 season were largely the same as those that would have been available if not for the hiatus, excepting a small number of opt-outs and injuries. There is also the possibility that different regions had varying levels of exposure to COVID-19, and that these differences could affect the home field advantage. To address this concern, we control for the difference between total cases, as of March 31, in the home compared to away team regions, which we then interact with the no-fans policy. The inclusion of these controls does not qualitatively affect our results.

It is also possible that, by chance, weaker teams tend to be at home during this period. Although this is unlikely due to the random process involved in scheduling, we control for three different measures of home versus away team strength: difference in season-to-date points, difference in points in the last 4 games, and difference in points at the end of the previous season.[20] Absent concerns with team-strength, there is still the concern that home field advantage fluctuates naturally throughout the season. The pressure on teams that are competing to stay above the relegation line or in the spots for European competitions could impact home field advantage. We address this by including (match) week by league fixed effects, which allow for differences in "seasonality" within each country.

There is also the possibility that the long hiatus of two or more months affected home and away teams differently. Within this channel are two subchannels, those being the long break itself and the differences in weather relative to when these games were originally scheduled. While there is some effect of these breaks on injuries (Ekstrand *et al.*, 2019) and the outcome of matches (Jamil *et al.*, 2020), there is no evidence that these breaks affect home and away teams differently. The primary impact of breaks on game performance is a decline in shot-to-goal conversion, but there is contradictory evidence that shot-to-goal conversion actually improved once play resumed from the COVID-induced hiatus (Cohen & Robinson, 2020). We abstract from potential differences in shot-conversion by analyzing expected goals in addition to realized goals, since expected goals reflect the quality and quantity of chances created regardless of the conversion rate. The second subchannel is that the delay in the season leads to different weather than in a normal season. The difference in average match-day temperature is evident in Figure 3.1, which plots the weather distribution for matches played with fans compared to without fans. There were two mitigation strategies undertaken by these four leagues

[20] Three points are earned for a won match, one point is earned for a drawn match, and no points are earned for a loss. For teams that are newly promoted to the top-division, we set their "previous season points" as the points for the team that barely avoided relegation in the previous season.

to limit the weather burden on players; teams were allowed five substitutions per game instead of the usual three, and there were two water breaks per game in addition to the halftime intermission. It is unlikely that these slight changes would affect home field advantage directly, but they could mitigate any differences in home field advantage driven by warmer weather than when the games were originally scheduled to be played. To account for any effects of weather, however, we include flexible, non-linear controls for average temperature.

The last potential confounding factor is that one of the other two channels of home field advantage — familiarity and travel — may have changed during this period. In stark contrast to many American sports leagues in 2020, games were played in the normal team stadia and were not moved to better mitigate the spread of COVID-19. Thus there was no difference in the familiarity effect as it is construed in the literature before and after the season was suspended. It is likely that travel was more onerous after the start of the pandemic; one might expect for away teams to perform worse relative to home teams due to more strenuous travel conditions.[21] This potential difference in travel would lead us to understate any negative effects that the removal of fans have on home field advantage. Regardless, we control for distance travelled by the away team to best isolate the impact of fans on home field advantage.

[21]If away teams are optimizing performance with respect to travel arrangements prior to the COVID-19 pandemic, then any additional restrictions placed upon those teams to avoid infection and outbreak could negatively affect the performance of away teams.

3.4 Empirical Framework

To recover a baseline estimate the effect of fans on home field advantage, we first estimate the following equation:

$$y_{ijmsct} = \alpha + \beta Post_{msct} + \epsilon_{ijmsct} \tag{3.1}$$

where the dependent variable, y_{ijmsct}, represents the home minus away realized or expected goals difference in matchup m of season s in league c and week t between home team i and away team j. The intercept α is the home field advantage under normal circumstances, capturing the total effect of all three factors (fans, familiarity, and travel), and $Post_{msct}$ is an indicator variable equal to 1 if matchup m occurs after the implementation of a no-fans policy. The coefficient of interest β represents the change in home field advantage once these no-fans policies are implemented. Standard errors ϵ_{ijmsct} are clustered at the matchup-by-season level to allow for correlation in the error term between observations involving the same two teams in the same season.[22]

The introduction of the no-fans policy induced by COVID-19 could coincide with changes in other factors that impact home field advantage. To mitigate any concerns of omitted variable bias, we estimate the following preferred specification:

$$y_{ijmsct} = \alpha + \beta Post_{msct} + \gamma X^1_{ijmsct} + \eta X^2_{ijmsct} + \lambda_{ct} + \epsilon_{ijmsct} \tag{3.2}$$

where we control for three sources of potential confounding factors. First, we address the concern that the introduction of the no-fans policy coincides with natural, within

[22]Note that for any matchup between two teams, there are two games played per season, one at each stadium. Thus, the home-away status in one game is perfectly negatively correlated with the home-away status of the same two teams in the reverse fixture.

season variation in home field advantage by including (match) week by league fixed effects, λ_{ct}. Second, we address the role of travel and form on home field advantage in X^1_{ijmsct}, which includes distance, distance squared, point difference entering the match, point difference in last 4 games, and point difference in the previous season.[23] Lastly, we account for two channels that are correlated with the no-fans policy — higher average temperatures and heterogeneous effects of COVID-19 by location — in X^2_{ijmsct}. The effect of temperature on home field advantage is likely to be non-linear and heterogeneous by league, due to differences in the typical climate faced. We allow for a flexible relationship between temperature and home field advantage by binning temperature at 5 degree Celsius intervals and allowing for a separate effect of temperature on home field advantage by league.[24] Lastly, we control for the heterogeneous effects of COVID-19 by location by controlling for the difference in team region COVID-19 cases, as of March 31, interacted with $Post_{msct}$.[25]

In addition to the above analysis of goal difference, we use an ordered logit model to investigate the effect of fans on match outcomes. For this specification, y_{msct}, takes on 3 values based on if the home-team won ($y_{msct} = 2$), drew ($y_{msct} = 1$), or lost ($y_{msct} = 0$) the match. We estimate the effect of the no-fans policy with the same controls shown in equation (3.2) and graphically show the changes in predicted probability of match outcomes.

[23] Since we are controlling for previous season point difference, the first season in our data, 2009-2010, is dropped. For teams that are promoted, the previous season points value entered is the lowest points value for a team that stayed up in the previous season. Our results are quantitatively robust to the exclusion of this variable and estimation on the full sample.

[24] This approach of binning is similar to what has been used when estimating the effect of temperature on mortality Deschênes & Greenstone (2011). In this context, we employ 6 temperature bins $\{< 0, 0-5, 5-10, 10-15, 15-20, > 20\}$ that are interacted with league indicator variables to account for potential non-linearity and heterogeneous effects by league.

[25] Our results are quantitatively robust to the use of either new cases on March 31 or total cases divided by population instead of overall cases.

3.5 Results

Column 1 of Table 3.4 shows that raw home field advantage decreased by 0.213 goals per game from a baseline of a 0.387 goals per game advantage for the home team.[26] This represents a decrease of 55%. However, some fraction of this decrease is driven by factors other than the no-fans policy. Column 4 shows that the coefficient on treatment drops by approximately 18% but remains statistically significant at the 10% level. These changes in estimates and inference highlight the importance of other factors, such as weather, in establishing the effect of fans on home field advantage using the COVID-19 natural experiment.

Column 1 of Table 3.5 shows that raw home field advantage, as measured by expected goals instead of realized goals, decreased by 64% from a 0.307 expected goal advantage for the home team to just 0.110 expected goals. Although the magnitude of the decrease is smaller than realized goals in absolute terms (0.197 xG as opposed to 0.213 G), it represents a larger fraction of the initial home field advantage (64% as opposed to 55%) because the initial home field advantage is smaller as measured by expected goals than realized goals. Column 4 of Table 3.5 shows that, although the effect of no fans on home field advantage, as measured by expected goals, decreases in magnitude when other controls are included, the effect remains statistically significant at the 1 percent level. This is in stark contrast to realized goals, in which the effect is only statistically significant at the 10 percent level when controls are included.[27]

In addition to the primary specification, we also estimate the effect of no fan policies on total goals per game. There is a small, statistically insignificant effect on total goals

[26]Results, shown in Table A.3.3 are qualitatively similar using only the restricted sample for which expected goals data are available.

[27]Realized goals are the result of a noisier process than expected goals. Each shot can only take on the

per game, shown in Table A.3.1; even the larger estimated coefficient would only indicate a 3.5% increase in total goals per game. Therefore, the shift in goal difference is derived from approximately equal parts of fewer goals for the home team and more goals for the away team. The same is true for total expected goals, in which the estimated coefficients are positive but statistically insignificant and represent a tiny increase relative to the baseline.

These shifts in goal difference due to the lack of fans manifest in fewer wins for the home teams. Figure 3.2 shows the simple shift in probability mass from home wins to home losses when fans are unable to attend games. Figure 3.3 shows the change in predicted probability of match outcomes from estimating an ordered logit. Importantly, it highlights that the lack of fans led to fewer home wins and more home losses, but the probability of a draw is unaffected, suggesting that fans are symmetrically pivotal: fans are approximately as likely to shift a result from a draw to a home win as they are from a home loss to a draw. Table A.3.2 presents qualitative similar results using a linear probability model instead of an ordered logit. We estimate a decrease in the probability of a home win and a precise null effect on the probability of a draw for the home team. Thus, there is a corresponding increase in probability for the omitted category (home loss). Approximately 5.4 percentage points are shifted from the probability of winning to the probability of losing.

In summary, we find that European soccer matches after the hiatus in the 2019-2020 season experienced a sharp decrease in the home field advantage across three key metrics. First, the actual difference between home and away goals decreased by 55 percent relative to the baseline. However, this decrease is only statistically significant at the 10 percent

values of zero or one realized goals, leading to larger error terms, whereas shots may relate to a value of expected goals anywhere between zero and one. This relates to the fact that expected goals are a better predictor of future performance than actual goals, as is discussed in Section 2.3.

level when controlling for other factors that impact home field advantage that could be correlated with the no-fans policy. Second, this drop in actual goals is representative of changes in the home field advantage in terms of chance creation, as measured by expected goals, which decreased by approximately the same amount. Expected goals are not subject to as much noise as realized goals, so this decrease is a more reliable gauge of the true effect of fans on home field advantage.[28] Lastly, these changes affect outcomes, as is shown by decreases in the probability of a home win and increases in the probability of a home loss.

3.6 Conclusion

Using exogenous variation in attendance due to the COVID-19 pandemic, we find that home field advantage decreased by 64% as measured by expected goals when fans were not allowed in stadia. This result is robust to the inclusion of various fixed effects and controls for many potential confounding factors, including weather, team quality, form, and COVID-19 cases by region. Our estimates also suggest that the home field advantage dropped from 0.387 realized goals to 0.174 goals when no fan policies were implemented to reduce the spread of the virus, but the estimated coefficient is only statistically significant at the 10 percent level because of the additional noise in the data generating process. These changes in home field advantage manifest in fewer home wins and more home losses when fans are not allowed in stadia. Using both an ordered logit regression and a linear probability model, we find that the lack of fans was important for match outcomes, suggesting that fans are themselves pivotal in how they affect sports.

[28]The effect of no-fans on expected goal difference is significant at the 1 percent level when controls are concluded, compared to the 10 percent level for realized goal difference.

One limitation of this paper is that we measure the effect of going from the "standard" number of fans down to zero. However, this standard level of fan attendance varies across leagues, teams, and even games, so the effect of marginal fans may vary in different contexts. Similarly, the impact of fans may be different for different types of competitions. Pollard & Pollard (2005) suggests that home advantage is smaller in national cup competitions, like the FA Cup in England, but larger in continental club competitions, like the UEFA Champions League. These different types of games also experience different levels of travel and fan attendance and may provide future evidence for the marginal effects of fans, as opposed to the total effect measured in this paper. Going forward, researchers may exploit plausibly exogenous variation in fan attendance for the 2020-2021 seasons, in which fan attendance was positive but limited to prevent the spread of the disease. The marginal effect of fans is a promising area of future research with real-world applications regarding optimal stadium capacity and ticket-pricing schema as soccer teams try to optimize on-field performance in addition to match-day revenue.

Another interesting note from the results is that the baseline home field advantage is approximately 24% smaller when measured by expected goals as opposed to realized goals. This may suggest that home-away differences in finishing (the act of converting scoring chances to goals) plays a significant role in the existing home field advantage. Future research may attempt to answer this question more directly, and analysis thereof may further explain and decompose home field advantage as it exists in soccer and other sports.

Tables

Table 3.1: Covid Summary by League

	Bundesliga	La Liga	Premier league	Serie A
Last Game before Pause	March 11	March 10	March 9	March 9
First Game after Pause	May 17	June 11	June 17	June 20
Days between Games	66	93	100	103
Games Cancelled	0	0	0	0
On-Field Changes	<——— 5 substitutes, Water Breaks ———>			
Away Travel	Multiple buses	Travel organized in-house	Fly, bus, or personal vehicles	Multiple buses
Fans Allowed?	No	No	No	No

Notes: This table includes some basic information on how COVID-19 impacted each of the leagues. All of the leagues stopped at about the same time, but restarts were staggered. No games were cancelled or postponed after the restart in any of the four major leagues.

Table 3.2: Summary Statistics for All Leagues

	Mean	Standard Error	Minimum	Maximum
Attendance (thousands)	32.52	18.74	0.01	259.30
Stadium Distance (km)	339.05	260.14	0.00	2259.93
Home Team Goals	1.57	1.32	0.00	10.00
Home Team Expected Goals	1.48	0.79	0.00	5.80
Away Team Goals	1.19	1.17	0.00	9.00
Away Team Expected Goals	1.19	0.71	0.00	5.30
Home-Away Goals Difference	0.38	1.83	-9.00	8.00
Home-Away Expected Goals Difference	0.29	1.15	-5.10	5.50

Notes: This table shows the summary statistics for all of the leagues. An observation is a game and stadium distance for each game is the distance between the two home stadia. Expected goals are only calculated for the 2017-2018 season onwards, but all other variables start in the 2009-2010 season.

Table 3.3: Summary Statistics by League

	Bundesliga	La Liga	Premier	Serie A
Attendance (thousands)	42.98	27.96	37.29	23.72
	(17.66)	(20.64)	(16.03)	(13.93)
Stadium Distance (km)	293.74	487.62	187.76	378.25
	(146.78)	(341.51)	(110.38)	(251.45)
Home Team Goals	1.64	1.59	1.57	1.52
	(1.36)	(1.36)	(1.32)	(1.25)
Home Team Expected Goals	1.59	1.43	1.43	1.48
	(0.82)	(0.75)	(0.79)	(0.80)
Away Team Goals	1.30	1.13	1.18	1.18
	(1.23)	(1.15)	(1.16)	(1.13)
Away Team Expected Goals	1.32	1.06	1.17	1.23
	(0.76)	(0.63)	(0.72)	(0.72)
Home-Away Goals Difference	0.34	0.46	0.38	0.34
	(1.94)	(1.85)	(1.84)	(1.69)
Home-Away Expected Goals Difference	0.28	0.37	0.26	0.25
	(1.23)	(0.98)	(1.21)	(1.19)

Notes: This table shows the differences between leagues across some key dimensions. Each cell represents the mean for that league, while the standard deviation is shown in parentheses. When calculating these statistics, an observation is a game and stadium distance for each game is the distance between the two home stadia. Expected goals are only calculated for the 2017-2018 season onwards, but all other variables start in the 2009-2010 season.

Table 3.4: Impact of No Fans on Home Field Advantage - Realized Goals

	(1)	(2)	(3)	(4)
	G_{H-A}	G_{H-A}	G_{H-A}	G_{H-A}
Baseline	0.387***	0.387***	0.336***	0.342***
	[0.013]	[0.013]	[0.029]	[0.032]
No Fans	-0.213**	-0.233**	-0.188**	-0.175*
	[0.100]	[0.103]	[0.088]	[0.105]
Points Difference			0.020***	0.020***
			[0.001]	[0.001]
Diff in Recent Points (4 Games)			0.019***	0.018***
			[0.004]	[0.005]
Diff in Prev. Season Points			0.022***	0.022***
			[0.001]	[0.001]
No Fans x Total Cases			-0.002	-0.002
			[0.009]	[0.008]
Controls	No	No	Yes	Yes
Week by League FE	No	Yes	No	Yes
Weather Controls	No	No	No	Yes
Observations	15,906	15,906	14,460	14,460
Games without Fans	408	408	408	408
Clusters	7,953	7,953	7,230	7,230

Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Notes: This table shows the change in home minus away goals when games are played with no fans (behind closed doors). This analysis uses data for all seasons from 2009-2020 in the following leagues: Bundesliga, Premier League, La Liga, and Serie A. The first row represents the estimated baseline home field advantage in terms of goals. The second row shows the estimated effect of the no-fans policy on home minus away goals. Each column shows a separate specification. The first column has no controls. The second column controls for league specific seasonality in home field advantage with (match) week by league fixed effects. The third column includes controls for team form and quality, distance traveled, and differences in COVID-19 cases by region interacted with the no-fans policy. The last column combines the previous two columns by including the controls of column 3 along with the fixed effects from column 2. In addition, we control for weather changes with 5 degree Celsius average temperature bin, $\{< 5, 0-5, 5-10, 10-15, 15-20, > 20\}$, by league fixed effects. Standard errors in brackets are clustered at the matchup by season level.

Table 3.5: Impact of No Fans on Home Field Advantage - Expected Goals

	(1) xG_{H-A}	(2) xG_{H-A}	(3) xG_{H-A}	(4) xG_{H-A}
Baseline	0.307***	0.308***	0.271***	0.285***
	[0.015]	[0.015]	[0.031]	[0.034]
No Fans	-0.197***	-0.204***	-0.184***	-0.181**
	[0.065]	[0.074]	[0.057]	[0.086]
Points Difference			0.010***	0.010***
			[0.002]	[0.002]
Diff in Recent Points (4 Games)			0.022***	0.024***
			[0.005]	[0.005]
Diff in Prev. Season Points			0.017***	0.016***
			[0.001]	[0.001]
No Fans x Total Cases			0.000	0.000
			[0.005]	[0.005]
Controls	No	No	Yes	Yes
Week by League FE	No	Yes	No	Yes
Weather Controls	No	No	No	Yes
Observations	4,336	4,336	4,336	4,336
Games without Fans	408	408	408	408
Clusters	2,169	2,169	2,169	2,169

Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Notes: This table shows the change in home minus away expected goals when games are played with no fans (behind closed doors). This analysis uses data for all seasons from 2017-2020 in the following leagues: Bundesliga, Premier League, La Liga, and Serie A. The first row represents the estimated baseline home field advantage in terms of expected goals. The second row shows the estimated effect of the no-fans policy on home minus away expected goals. Each column shows a separate specification. The first column has no controls. The second column controls for league specific seasonality in home field advantage with (match) week by league fixed effects. The third column includes controls for team form and quality, distance traveled, and differences in COVID-19 cases by region interacted with the no-fans policy. The last column combines the previous two columns by including the controls of column 3 along with the fixed effects from column 2. In addition, we control for weather changes with 5 degree Celsius average temperature bin, $\{< 5, 0-5, 5-10, 10-15, 15-20, > 20\}$, by league fixed effects. Standard errors in brackets are clustered at the matchup by season level.

Figures

Figure 3.1: Weather Distribution

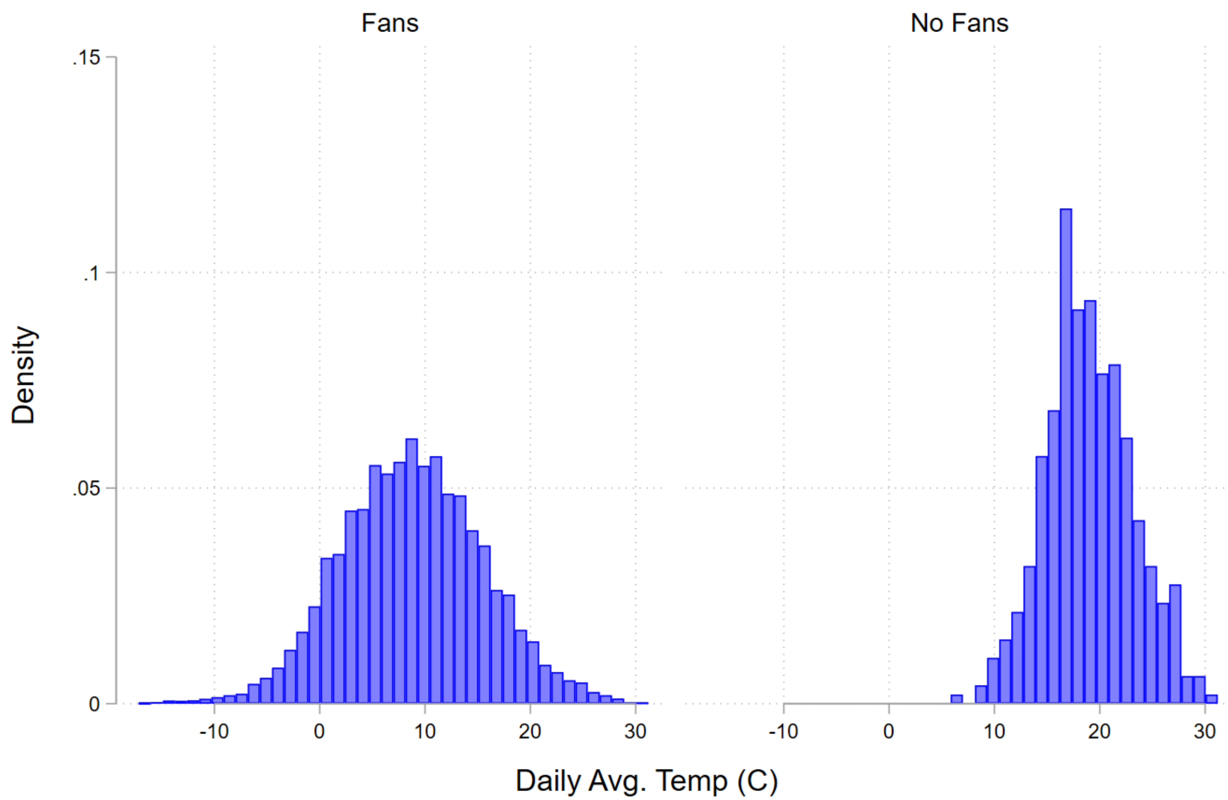

Note: This figure shows the distribution of average daily temperatures for matches played when there are no fans compared to when there are fans. There is a shift in the distribution to higher temperatures, which reflects that matches were pushed into the summer.

Figure 3.2: Change in Match Outcomes with No Fans

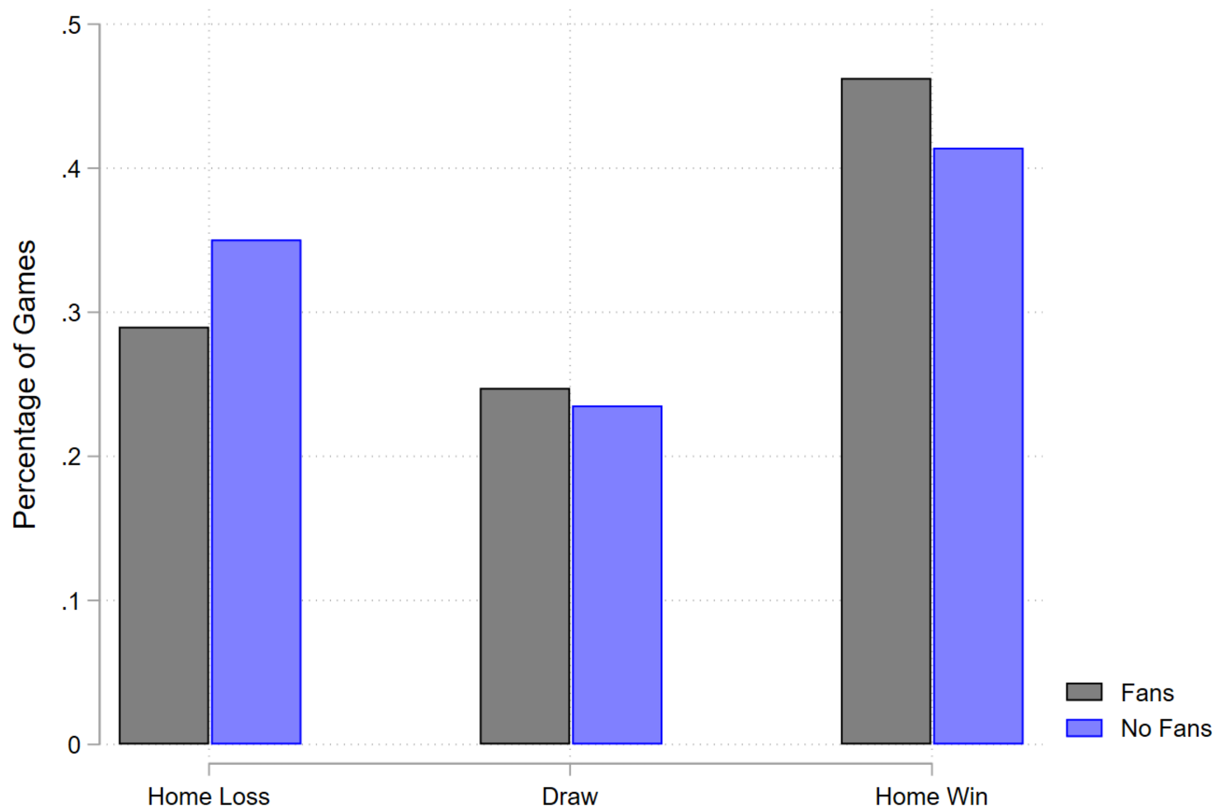

Note: This figure shows the distribution of match outcomes when there are no fans compared to when there are fans. There is little change in the percentage of games that end as a draw, but there are large differences for wins and losses.

Figure 3.3: Effect of No Fans on Match Outcomes

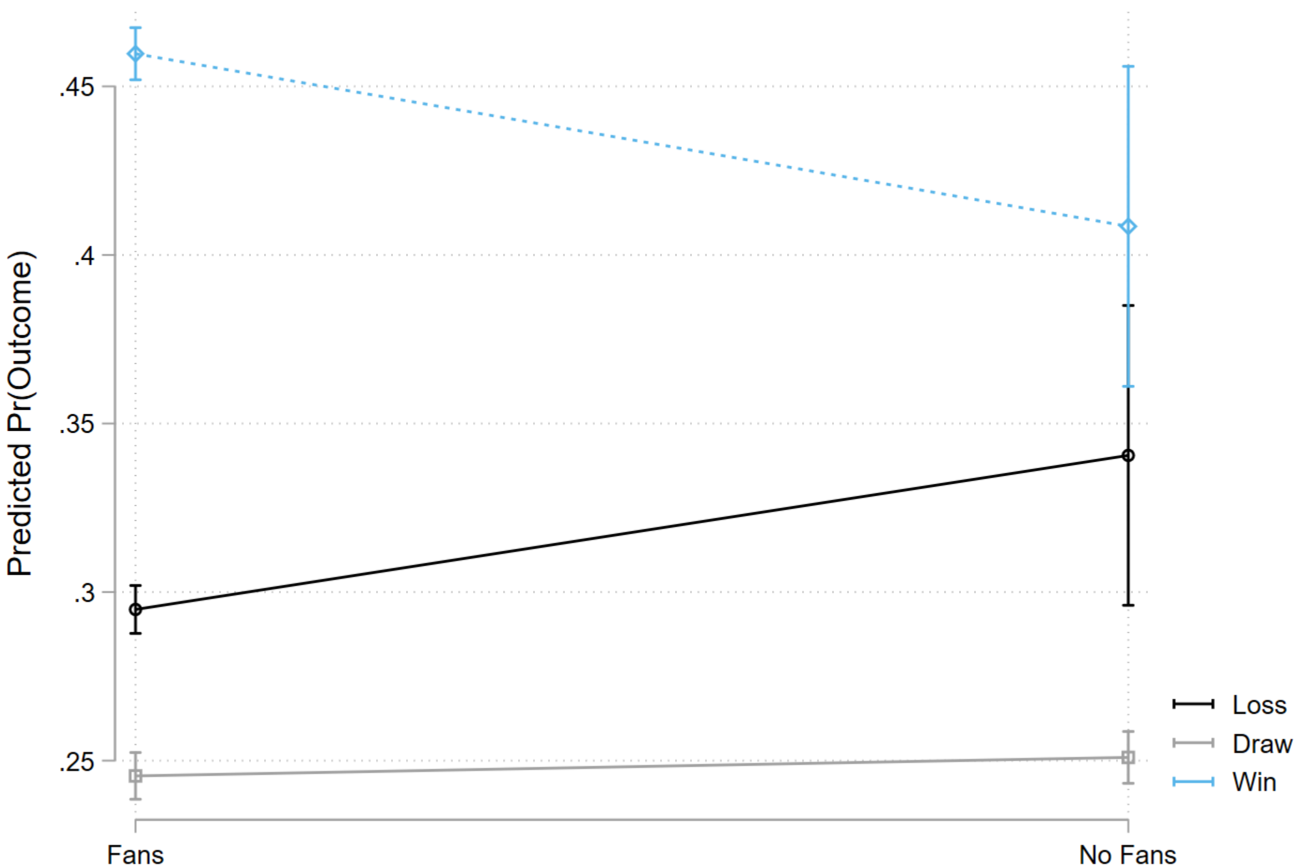

Note: This figure shows the predicted probability of win, draw, and loss before and after the no-fans policy. Average predicted probabilities are calculated after fitting an ordered logistic model for match outcomes that matches the specification estimated in column 4 of Tables 3.4 and 3.5. 95 percent confidence intervals are represented in the bands.

Appendix A

Additional Tables and Figures

A.1 Chapter 1

Table A.1: Staggered Adoption Differences in Differences - Contemporaneous Effect

	Rural	Urban	Rural 11+	Urban 11+
	prob(att)	prob(att)	prob(att)	prob(att)
Contemporaneous Treatment	0.0273	-0.0271	0.0342	-0.0170
(observed after law passed)	(0.0261)	(0.0150)	(0.0278)	(0.0139)
Full Treatment	0.0466*	0.0254	0.0517*	0.0309
($age < 7$ when law passed)	(0.0202)	(0.0247)	(0.0204)	(0.0244)
Partial Treatment	0.0111	0.0263	0.0135	0.0331*
($11 > age > 6$ when law passed)	(0.0107)	(0.0138)	(0.0143)	(0.0157)
N	17,610,766	5,704,146	6,544,064	2,045,437

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Standard errors clustered at the state level in parentheses. Critical values are obtained from the t-distribution with 26 degrees of freedom. Column 1 shows the results from Equation 1 estimated on rural individuals between the ages of 5 and 14. Column 2 shows the results from Equation 1 estimated on urban individuals between the ages of 5 and 14. Column 3 shows the results from Equation 1 estimated on rural individuals between the ages of 10 and 14. Column 4 shows the results from Equation 1 estimated on urban individuals between the ages of 10 and 14. All regressions include state-by-age fixed effects and Census year-by-age fixed effects, as well as a constant term. No other controls are included.

Table A.2: Staggered Adoption Differences in Differences - Contemporaneous Effect, Low SES Only

	Rural prob(att)	Urban prob(att)	Rural 11+ prob(att)	Urban 11+ prob(att)
Contemporaneous Treatment	0.0271	-0.0299	0.0346	-0.0188
(observed after law passed)	(0.0263)	(0.0153)	(0.0281)	(0.0149)
Full Treatment	0.0486*	0.0268	0.0543*	0.0332
($age < 7$ when law passed)	(0.0206)	(0.0257)	(0.0210)	(0.0250)
Partial Treatment	0.0122	0.0284	0.0152	0.0364*
($11 > age > 6$ when law passed)	(0.0110)	(0.0152)	(0.0145)	(0.0175)
N	16,222,381	4,441,612	6,024,103	1,572,870

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Standard errors clustered at the state level in parentheses. Critical values are obtained from the t-distribution with 26 degrees of freedom. Column 1 shows the results from Equation 1 estimated on low-SES rural individuals between the ages of 5 and 14. Column 2 shows the results from Equation 1 estimated on urban individuals between the ages of 5 and 14. Column 3 shows the results from Equation 1 estimated on rural individuals between the ages of 10 and 14. Column 4 shows the results from Equation 1 estimated on urban individuals between the ages of 10 and 14. All regressions include state-by-age fixed effects and Census year-by-age fixed effects, as well as a constant term. No other controls are included.

Table A.3: Staggered Adoption Differences in Differences - High SES Only

	Rural prob(att)	Urban prob(att)	Rural 10+ prob(att)	Urban 10+ prob(att)
Full Treatment	0.0571*	0.0018	0.0618**	0.0160
($age < 7$	(0.0209)	(0.0227)	(0.0220)	(0.0228)
when law passed)	[0.015,0.118]**	[-0.048,0.078]	[0.017,0.127]**	[-0.036,0.089]
Partial Treatment	0.0340	0.0025	0.0368	0.0153
($11 > age > 6$	(0.0208)	(0.0140)	(0.0215)	(0.0128)
when law passed)	[-0.008,0.091]	[-0.021,0.047]	[-0.006,0.097]	[-0.006,0.059]
N	6,618,741	4,949,965	3,103,894	2,275,181

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Standard errors clustered at the state level in parentheses. Critical values are obtained from the t-distribution with 26 degrees of freedom. Column 1 shows the results from Equation 3 estimated on high-SES rural individuals between the ages of 5 and 14. Column 2 shows the results from Equation 3 estimated on urban individuals between the ages of 5 and 14. Column 3 shows the results from Equation 3 estimated on rural individuals between the ages of 10 and 14. Column 4 shows the results from Equation 3 estimated on urban individuals between the ages of 10 and 14. All regressions include state-by-age fixed effects and Census year-by-age fixed effects, as well as a constant term. No other controls are included.

Table A.4: Staggered Adoption Differences in Differences - Low SES Only, South Excluded

	Rural prob(att)	Urban prob(att)	Rural 10+ prob(att)	Urban 10+ prob(att)
Full Treatment	0.0454*	-0.0257	0.0438*	-0.0117
($age < 7$	(0.0193)	(0.0178)	(0.0183)	(0.0227)
when law passed)	[0.005,0.115]*	[-0.060,0.030]	[0.005,0.105]*	[-0.062,0.058]
Partial Treatment	0.0133	-0.0254	0.0126	-0.0149
($11 > age > 6$	(0.0201)	(0.0134)	(0.0197)	(0.0126)
when law passed)	[-0.029,0.084]	[-0.060,0.010]	[-0.030,0.080]	[-0.048,0.021]
N	7,993,373	661,282	3,861,848	315,797

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Standard errors clustered at the state level in parentheses. Critical values are obtained from the t-distribution with 17 degrees of freedom. Column 1 shows the results from Equation 3 estimated on rural individuals between the ages of 5 and 14. Column 2 shows the results from Equation 3 estimated on urban individuals between the ages of 5 and 14. Column 3 shows the results from Equation 3 estimated on rural individuals between the ages of 10 and 14. Column 4 shows the results from Equation 3 estimated on urban individuals between the ages of 10 and 14. All regressions include state-by-age fixed effects and Census year-by-age fixed effects, as well as a constant term. No other controls are included.

Table A.5: Staggered Adoption Differences in Differences - High SES Only, South Excluded

	Rural prob(att)	Urban prob(att)	Rural 10+ prob(att)	Urban 10+ prob(att)
Full Treatment	0.0305	-0.0171	0.0336	0.0017
($age < 7$	(0.0187)	(0.0178)	(0.0178)	(0.0192)
when law passed)	[-0.004,0.093]	[-0.053,0.046]	[-0.001,0.091]	[-0.037,0.071]
Partial Treatment	0.0053	-0.0162	0.0057	-0.0001
($11 > age > 6$	(0.0167)	(0.0086)	(0.0165)	(0.0057)
when law passed)	[-0.024,0.065]	[-0.035,0.007	[-0.023,0.066]	[-0.011,0.016]
N	5,313,697	4,354,067	2,482,693	1,999,381

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Standard errors clustered at the state level in parentheses. Critical values are obtained from the t-distribution with 17 degrees of freedom. Column 1 shows the results from Equation 3 estimated on rural individuals between the ages of 5 and 14. Column 2 shows the results from Equation 3 estimated on urban individuals between the ages of 5 and 14. Column 3 shows the results from Equation 3 estimated on rural individuals between the ages of 10 and 14. Column 4 shows the results from Equation 3 estimated on urban individuals between the ages of 10 and 14. All regressions include state-by-age fixed effects and Census year-by-age fixed effects, as well as a constant term. No other controls are included.

Table A.6: Staggered Adoption Differences in Differences - Alternate Definition of Treatment

	Rural prob(att)	Urban prob(att)	Rural 10+ prob(att)	Urban 10+ prob(att)
Full Treatment	0.0721**	0.0003	0.0790**	0.0148
($age < 7$	(0.0202)	(0.0225)	(0.0232)	(0.0227)
when law passed)	[0.029,0.127]***	[-0.050,0.074]	[0.031,0.142]***	[-0.038,0.087]
Partial Treatment	0.0398	-0.0004	0.0533*	0.0187
($age = 7$ wlp)	(0.0201)	(0.0160)	(0.0227)	(0.0173)
Partial Treatment	0.0359	0.0024	0.0375	0.0155
($age = 8$ wlp)	(0.0213)	(0.0167)	(0.0240)	(0.0157)
Partial Treatment	0.0399	0.0044	0.0438	0.0133
($age = 9$ wlp)	(0.0212)	(0.0134)	(0.0220)	(0.0119)
Partial Treatment	0.0281	-0.0002	0.0316	0.0079
($age = 10$ wlp)	(0.0234)	(0.0136)	(0.0240)	(0.0117)
N	17,610,766	5,704,146	8,398,881	2,636,717

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Standard errors clustered at the state level in parentheses. Critical values are obtained from the t-distribution with 26 degrees of freedom. Column 1 shows the results from Equation 3 estimated on rural individuals between the ages of 5 and 14. Column 2 shows the results from Equation 3 estimated on urban individuals between the ages of 5 and 14. Column 3 shows the results from Equation 3 estimated on rural individuals between the ages of 10 and 14. Column 4 shows the results from Equation 3 estimated on urban individuals between the ages of 10 and 14. All regressions include state-by-age fixed effects and Census year-by-age fixed effects, as well as a constant term. No other controls are included.

Table A.7: Staggered Adoption Differences in Differences - State Natives Only

	Rural prob(att)	Urban prob(att)	Rural 10+ prob(att)	Urban 10+ prob(att)
Full Treatment	0.0617**	-0.0056	0.0660*	0.0070
($age < 7$	(0.0214)	(0.0220)	(0.0253)	(0.0228)
when law passed)	[0.018,0.123]**	[-0.051,0.070]	[0.015,0.142]**	[-0.042,0.085]
Partial Treatment	0.0293	0.0003	0.0330	0.0110
($11 > age > 6$	(0.0233)	(0.0133)	(0.0242)	(0.0127)
when law passed)	[-0.020,0.087]	[-0.022,0.044]	[-0.022,0.099]	[-0.010,0.055]
N	15,094,635	5,043,825	7,009,657	2,284,139

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Standard errors clustered at the state level in parentheses. Critical values are obtained from the t-distribution with 26 degrees of freedom. Column 1 shows the results from Equation 1 estimated on rural individuals between the ages of 5 and 14, excluding those individuals that did not live in the same state of their birth. Column 2 shows the results from Equation 1 estimated on urban individuals between the ages of 5 and 14. Column 3 shows the results from Equation 1 estimated on rural individuals between the ages of 10 and 14. Column 4 shows the results from Equation 1 estimated on urban individuals between the ages of 10 and 14. All regressions include state-by-age fixed effects and Census year-by-age fixed effects, as well as a constant term. No other controls are included.

A.2 Chapter 2

Table A.1: Rural-Urban Fertility Differences among Married Women

	(1) Refugee	(2) NRI	(3) Native	(4) Refugee	(5) NRI	(6) Native
Urban	-0.011**	-0.009***	-0.010***	-0.004	-0.004*	0.005*
	(0.004)	(0.002)	(0.001)	(0.005)	(0.002)	(0.002)
Sample	All	All	All	Married	Married	Married
Age Fixed Effects	X	X	X	X	X	X
Educ Fixed Effects	X	X	X	X	X	X
# Kids Fixed Effects	X	X	X	X	X	X
POB Fixed Effects	X	X	X	X	X	X
Clustered?	POB	POB	POB	POB	POB	POB

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Standard errors clustered at the ”Place of Birth” level in parentheses, where the “Place of birth” refers to either the province or foreign country or birth. Column 1 shows the rural-urban fertility differences among refugees after the inclusion of various fixed effects to control for some observable characteristics. Columns 2 and 3 show the same differences for non-refugee immigrants and native Canadians, respectively. Column 4 shows the rural-urban fertility differences among married refugee women. Columns 5 and 6 show the same differences for non-refugee immigrants and native Canadians, respectively. Estimated coefficients for married women only are closer to zero and statistically insignificant for refugees. This may suggest that marital behavior differs in rural and urban settings.

Table A.2: Rural-Urban Fertility Differences among Non-Migrants

	(1) Refugee	(2) NRI	(3) Native	(4) Refugee	(5) NRI	(6) Native
Urban	-0.011**	-0.009***	-0.010***	-0.011**	-0.009***	-0.010***
	(0.004)	(0.002)	(0.001)	(0.004)	(0.002)	(0.001)
Sample	All	All	All	Limited	Limited	Limited
Age Fixed Effects	X	X	X	X	X	X
Educ Fixed Effects	X	X	X	X	X	X
# Kids Fixed Effects	X	X	X	X	X	X
POB Fixed Effects	X	X	X	X	X	X
Clustered?	POB	POB	POB	POB	POB	POB

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Standard errors clustered at the ”Place of Birth” level in parentheses, where the “Place of birth” refers to either the province or foreign country or birth. “Limited” Sample here refers specifically to non-migrants. Column 1 shows the rural-urban fertility differences among refugees after the inclusion of various fixed effects to control for some observable characteristics. Columns 2 and 3 show the same differences for non-refugee immigrants and native Canadians, respectively. Column 4 shows the rural-urban fertility differences among refugee women that live in the same Census Metropolitan Area as five years ago. Columns 5 and 6 show the same differences for non-refugee immigrants and native Canadians, respectively. Estimated coefficients are quantitatively similar. Although this does not guarantee that refugees live where they were exogenously placed when they entered Canada, the fact that these coefficients are identical to the full sample suggests that geographic mobility does not play a role in fertility decisions.

A.3 Chapter 3

Figure A.3.1: Distance from Stadium to Station

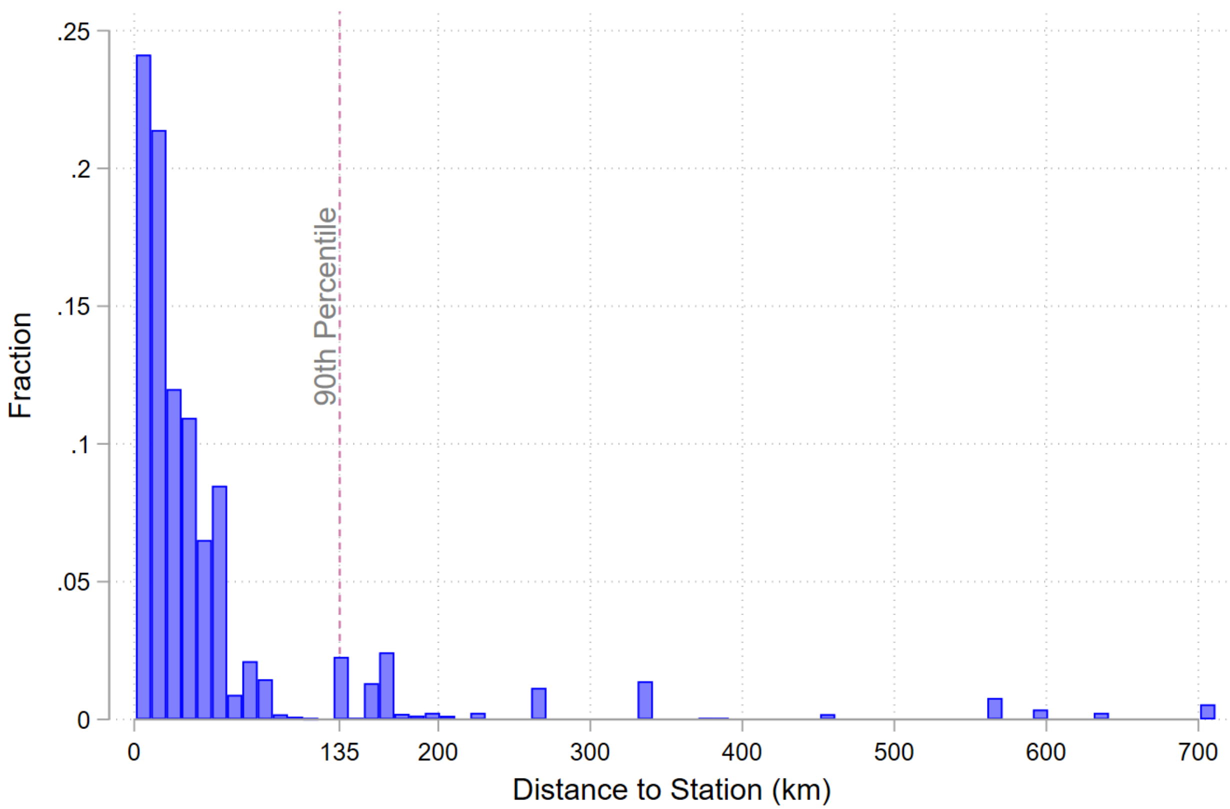

Note: This figure shows the distribution of distances between the stadium and stations. The gray line indicates the 90th percentile, which is 135 kilometers.

Table A.3.1: Impact of No Fans on Total Goals per Game

	(1)	(2)	(3)	(4)
	G_{total}	G_{total}	xG_{total}	xG_{total}
Baseline	2.766***	2.732***	2.661***	2.624***
	[0.014]	[0.034]	[0.016]	[0.036]
No Fans	0.099	-0.032	0.049	0.022
	[0.079]	[0.093]	[0.051]	[0.074]
Controls	No	Yes	No	Yes
Week by League FE	No	Yes	No	Yes
Observations	15,906	14,460	4,336	4,336
Games without Fans	408	408	408	408
Clusters	7,953	7,230	2,169	2,169

Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Notes: This table shows the change in goals and total goals when games are played with no fans (behind closed doors). This analysis uses data for all seasons from 2009-2020 for the first two columns and 2017-2020 for the last two columns in the following leagues: Bundesliga, Premier League, La Liga, and Serie A. The first row represents the estimated baseline average total goals or total expected goals. The second row shows the estimated effect of the no-fans policy on total goals or total expected goals. Each column shows a separate specification or dependent variable. The first and third columns have no controls. The second and fourth columns control for league specific seasonality in home field advantage with (match) week by league fixed effects. They also include controls for team form and quality, distance traveled, and differences in COVID-19 cases by region interacted with the no-fans policy. In addition, we control for weather changes with 5 degree Celsius average temperature bin, $\{< 5, 0-5, 5-10, 10-15, 15-20, > 20\}$, by league fixed effects. Standard errors in brackets are clustered at the matchup by season level.

Table A.3.2: Impact of No Fans on Home Field Advantage - Outcomes

	(1) Win	(2) Win	(3) Draw	(4) Draw
Baseline	0.463***	0.446***	0.247***	0.246***
	[0.004]	[0.009]	[0.003]	[0.008]
No Fans	-0.048*	-0.057**	-0.012	0.023
	[0.025]	[0.027]	[0.021]	[0.025]
Controls	No	Yes	No	Yes
Week by League FE	No	Yes	No	Yes
Observations	15,906	14,460	15,906	14,460
Games without Fans	408	408	408	408
Clusters	7,953	7,230	7,953	7,230

Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Notes: This table shows the change in the probability of a home team win or draw when games are played with no fans (behind closed doors). This analysis uses data for all seasons from 2009-2020 in the following leagues: Bundesliga, Premier League, La Liga, and Serie A. The first row represents the estimated baseline probability of a home team win or draw. The second row shows the estimated effect of the no-fans policy on the match outcomes. Each column shows a separate specification or dependent variable. The first and third columns have no controls. The second and fourth columns control for league specific seasonality in home field advantage with (match) week by league fixed effects. They also include controls for team form and quality, distance traveled, and differences in COVID-19 cases by region interacted with the no-fans policy. In addition, we control for weather changes with 5 degree Celsius average temperature bin, $\{< 5, 0-5, 5-10, 10-15, 15-20, > 20\}$, by league fixed effects. Standard errors in brackets are clustered at the matchup by season level.

Table A.3.3: Restricted Sample Realized Goals

	(1) G_{H-A}	(2) G_{H-A}	(3) G_{H-A}	(4) G_{H-A}
Baseline	0.329***	0.336***	0.292***	0.282***
	[0.026]	[0.026]	[0.052]	[0.058]
No Fans	-0.174*	-0.245**	-0.158*	-0.273*
	[0.104]	[0.118]	[0.091]	[0.142]
General Controls	No	No	Yes	Yes
Week by League FE	No	Yes	No	Yes
Weather Controls	No	No	No	Yes
Observations	4,336	4,336	4,336	4,336
Games without Fans	408	408	408	408
Clusters	2,169	2,169	2,169	2,169

Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Notes: This table shows the change in home minus away goals when games are played with no fans (behind closed doors). This analysis uses data for all seasons from 2017-2020 in the following leagues: Bundesliga, Premier League, La Liga, and Serie A. The first row represents the estimated baseline home field advantage in terms of goals. The second row shows the estimated effect of the no-fans policy on home minus away goals. Each column shows a separate specification. The first column has no controls. The second column controls for league specific seasonality in home field advantage with (match) week by league fixed effects. The third column includes controls for team form and quality, distance traveled, and differences in COVID-19 cases by region interacted with the no-fans policy. The last column combines the previous two columns by including the controls of column 3 along with the fixed effects from column 2. In addition, we control for weather changes with 5 degree Celsius average temperature bin, $\{< 5, 0-5, 5-10, 10-15, 15-20, > 20\}$, by league fixed effects. Standard errors in brackets are clustered at the matchup by season level.

Bibliography

Agnew, Gary A, & Carron, Albert V. 1994. Crowd effects and the home advantage. *International Journal of Sport Psychology.*

Baicker, Katherine, Mullainathan, Sendhil, & Schwartzstein, Joshua. 2015. Behavioral hazard in health insurance. *The Quarterly Journal of Economics*, **130**(4), 1623–1667.

Barnett, V, & Hilditch, S. 1993. The effect of an artificial pitch surface on home team performance in football (soccer). *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, **156**(1), 39–50.

Becker, Gary S. 1960. An Economic Analysis of Fertility, Demographic and economic change in developed countries: a conference of the Universities. *National Bureau Commitee for Economic Research*, **209**.

Belchior, Carlos Alberto. 2020. Fans and Match Results: Evidence From a Natural Experiment in Brazil. *Journal of Sports Economics*, 1527002520930812.

Benz, Luke S, & Lopez, Michael J. 2021. Estimating the change in soccer's home advantage during the Covid-19 pandemic using bivariate Poisson regression. *AStA Advances in Statistical Analysis*, 1–28.

Bertrand, Marianne, Duflo, Esther, & Mullainathan, Sendhil. 2004. How much should we trust differences-in-differences estimates? *The Quarterly Journal of Economics*, **119**(1), 249–275.

Black, Dan A, Kolesnikova, Natalia, Sanders, Seth G, & Taylor, Lowell J. 2013. Are children "normal"? *The review of economics and statistics*, **95**(1), 21–33.

Black, Sandra E, Devereux, Paul J, & Salvanes, Kjell G. 2008. Staying in the classroom and out of the maternity ward? The effect of compulsory schooling laws on teenage births. *The Economic Journal*, **118**(530), 1025–1054.

Bryson, Alex, Dolton, Peter, Reade, J James, Schreyer, Dominik, & Singleton, Carl. 2021. Causal effects of an absent crowd on performances and refereeing decisions during Covid-19. *Economics Letters*, **198**, 109664.

Cameron, A Colin, Gelbach, Jonah B, & Miller, Douglas L. 2008. Bootstrap-based improvements for inference with clustered errors. *The Review of Economics and Statistics*, **90**(3), 414–427.

Cao, Kang Hua, & Birchenall, Javier A. 2013. Agricultural productivity, structural change, and economic growth in post-reform China. *Journal of Development Economics*, **104**, 165–180.

Card, David. 1990. The impact of the Mariel boatlift on the Miami labor market. *ILR Review*, **43**(2), 245–257.

Carleton, Tamma, Cornetet, Jules, Huybers, Peter, Meng, Kyle C, & Proctor, Jonathan. 2021. Global evidence for ultraviolet radiation decreasing COVID-19 growth rates. *Proceedings of the National Academy of Sciences*, **118**(1).

Carter, Andrew V, Schnepel, Kevin T, & Steigerwald, Douglas G. 2017. Asymptotic behavior of at-test robust to cluster heterogeneity. *Review of Economics and Statistics*, **99**(4), 698–709.

Census Bureau, US. 1975. *Historical Statistics of the United States: Colonial Times to 1970*. Bicentennial edn.

Choudhry, Niteesh K, Avorn, Jerry, Glynn, Robert J, Antman, Elliott M, Schneeweiss, Sebastian, Toscano, Michele, Reisman, Lonny, Fernandes, Joaquim, Spettell, Claire, Lee, Joy L, *et al.* 2011. Full coverage for preventive medications after myocardial infarction. *New England Journal of Medicine*, **365**(22), 2088–2097.

Clark, William AV. 2012. Do women delay family formation in expensive housing markets? *Demographic research*, **27**(1), 1.

Clarke, Stephen R, & Norman, John M. 1995. Home ground advantage of individual clubs in English soccer. *Journal of the Royal Statistical Society: Series D (The Statistician)*, **44**(4), 509–521.

Cohen, Ben, & Robinson, Joshua. 2020. The world's best athletes are now better at shooting; NBA stars in the bubble and european soccer players in empty stadiums have one thing in common: They have improved at shooting. *Wall Street Journal (Online)*.

Cubberley, Ellwood Patterson. 1919. *Public education in the United States: A study and interpretation of American educational history; an introductory textbook dealing with the larger problems of present-day education in the light of their historical development.* Houghton Mifflin.

Denison, Edward F. 1985. Trends in Economic Growth, 1929-1982. *Brookings, Washington DC*.

Denning, Jeffrey T. 2017. College on the cheap: Consequences of community college tuition reductions. *American Economic Journal: Economic Policy*, **9**(2), 155–88.

Deschênes, Olivier, & Greenstone, Michael. 2011. Climate change, mortality, and adaptation: Evidence from annual fluctuations in weather in the US. *American Economic Journal: Applied Economics*, **3**(4), 152–85.

Dowie, Jack. 1982. Why Spain should win the world cup. *New Scientist*, **94**(10), 693–695.

Duflo, Esther. 2001. Schooling and labor market consequences of school construction in Indonesia: Evidence from an unusual policy experiment. *American Economic Review*, **91**(4), 795–813.

Ekstrand, Jan, Spreco, Armin, & Davison, Michael. 2019. Elite football teams that do not have a winter break lose on average 303 player-days more per season to injuries than those teams that do: a comparison among 35 professional European teams. *British journal of sports medicine*, bjsports–2018.

Emerick, Kyle. 2018. Agricultural productivity and the sectoral reallocation of labor in rural India. *Journal of Development Economics*, **135**, 488–503.

Endrich, Marek, & Gesche, Tobias. 2020. Home-bias in referee decisions: Evidence from "Ghost Matches" during the Covid19-Pandemic. *Economics Letters*, **197**, 109621.

Ferraresi, Massimiliano, Gucciardi, Gianluca, *et al.* 2020. Team performance and audience: experimental evidence from the football sector.

Fischer, Kai, & Haucap, Justus. 2021. Does crowd support drive the home advantage in professional football? Evidence from German ghost games during the COVID-19 pandemic. *Journal of Sports Economics*, **22**(8), 982–1008.

Foged, Mette, & Peri, Giovanni. 2016. Immigrants' effect on native workers: New analysis on longitudinal data. *American Economic Journal: Applied Economics*, **8**(2), 1–34.

Go, Sun. 2009. Free Schools in America, 1850-1870: Who Voted for them, who got them, and who paid for them.

Go, Sun. 2015. Free Elementary Schooling in Connecticut, 1866-1880. *Economic History*, **59**(4), 49–75.

Go, Sun, & Lindert, Peter. 2010. The uneven rise of American public schools to 1850. *The Journal of Economic History*, **70**(1), 1–26.

Goldin, Claudia. 1998. America's graduation from high school: The evolution and spread of secondary schooling in the twentieth century. *The Journal of Economic History*, **58**(2), 345–374.

Goldin, Claudia, & Katz, Lawrence. 2008. The race between technology and education. *Cambridge, MA: Harvard*.

Goldin, Claudia, & Katz, Lawrence F. 2000. Education and income in the early twentieth century: Evidence from the prairies. *The Journal of Economic History*, **60**(3), 782–818.

Goumas, Chris. 2014. Home advantage in Australian soccer. *Journal of Science and Medicine in Sport*, **17**(1), 119–123.

Gross, Tal, Layton, Timothy, & Prinz, Daniel. 2020. *The Liquidity Sensitivity of Healthcare Consumption: Evidence from Social Security Payments*. Tech. rept. National Bureau of Economic Research.

Harris, Brian L, Stergachis, Andy, & Ried, L Douglas. 1990. The effect of drug co-payments on utilization and cost of pharmaceuticals in a health maintenance organization. *Medical Care*, 907–917.

Jamil, Mikael, McErlain-Naylor, Stuart A, & Beato, Marco. 2020. Investigating the impact of the mid-season winter break on technical performance levels across European football–Does a break in play affect team momentum? *International Journal of Performance Analysis in Sport*, 1–14.

Klein Tank, AMG, Wijngaard, JB, Können, GP, Böhm, Reinhart, Demarée, Gaston, Gocheva, A, Mileta, M, Pashiardis, S, Hejkrlik, L, Kern-Hansen, C, *et al.* 2002. Daily dataset of 20th-century surface air temperature and precipitation series for the European Climate Assessment. *International Journal of Climatology: A Journal of the Royal Meteorological Society*, **22**(12), 1441–1453.

Kulu, Hill. 2013. Why do fertility levels vary between urban and rural areas? *Regional studies*, **47**(6), 895–912.

Kulu, Hill, & Washbrook, Elizabeth. 2014. Residential context, migration and fertility in a modern urban society. *Advances in Life Course Research*, **21**, 168–182.

Kuznets, Simon. 1974. Rural-urban differences in fertility: An international comparison. *Proceedings of the American Philosophical Society*, **118**(1), 1–29.

Landes, William M, & Solmon, Lewis C. 1972. Compulsory schooling legislation: An economic analysis of law and social change in the nineteenth century. *The Journal of*

Economic History, **32**(1), 54–91.

Lerch, Mathias. 2018. Fertility decline in urban and rural areas of developing countries. *Population and Development Review*, 1–20.

Lindert, Peter H. 2004. *Growing public: Volume 1, the story: Social spending and economic growth since the eighteenth century*. Vol. 1. Cambridge University Press.

Loughead, Todd M, Carron, Albert V, Bray, Steven R, & Kim, Arvin J. 2003. Facility familiarity and the home advantage in professional sports. *International Journal of Sport and Exercise Psychology*, **1**(3), 264–274.

MacKinnon, James G, & Webb, Matthew D. 2017. *Pitfalls when estimating treatment effects using clustered data*. Tech. rept. Queen's Economics Department Working Paper.

Marcotte, Dave E. 2019. The returns to education at community colleges: new evidence from the Education Longitudinal Survey. *Education Finance and Policy*, **14**(4), 523–547.

Marcotte, Dave E, Bailey, Thomas, Borkoski, Carey, & Kienzl, Greg S. 2005. The returns of a community college education: Evidence from the National Education Longitudinal Survey. *Educational Evaluation and Policy Analysis*, **27**(2), 157–175.

Mead, Arthur Raymond. 1918. *The development of free schools In the United States as Illustrated by Connecticut and Michigan*. Teachers college, Columbia university.

Nichols, Mark W. 2014. The impact of visiting team travel on game outcome and biases in NFL betting markets. *Journal of Sports Economics*, **15**(1), 78–96.

Oberhofer, Harald, Philippovich, Tassilo, & Winner, Hannes. 2010. Distance matters in away games: Evidence from the German football league. *Journal of Economic Psychology*, **31**(2), 200–211.

Pollard, Richard. 2002. Evidence of a reduced home advantage when a team moves to a new stadium. *Journal of Sports Sciences*, **20**(12), 969–973.

Pollard, Richard, & Pollard, Gregory. 2005. Home advantage in soccer: A review of its existence and causes.

Ponzo, Michela, & Scoppa, Vincenzo. 2018. Does the home advantage depend on crowd support? Evidence from same-stadium derbies. *Journal of Sports Economics*, **19**(4), 562–582.

Rathke, Alex. 2017. An examination of expected goals and shot efficiency in soccer. *Journal of Human Sport and Exercise*, **12**(2), 514–529.

Reade, J James, Schreyer, Dominik, & Singleton, Carl. 2021. Eliminating supportive crowds reduces referee bias. *Economic Inquiry*.

Roback, Jennifer. 1982. Wages, rents, and the quality of life. *Journal of political Economy*, **90**(6), 1257–1278.

Ruggles, Steven, Flood, Sarah, Goeken, Ronald, Grover, Josiah, Meyer, Erin, Pacas, Jose, & Sobek, Matthew. 2019. *IPUMS USA: Version 9.0 [dataset]*.

Scoppa, Vincenzo. 2021. Social pressure in the stadiums: Do agents change behavior without crowd support? *Journal of economic psychology*, **82**, 102344.

Simon, Curtis J, & Tamura, Robert. 2009. Do higher rents discourage fertility? Evidence from US cities, 1940–2000. *Regional Science and Urban Economics*, **39**(1), 33–42.

Swift, Fletcher Harper. 1911. *A history of public permanent common school funds in the United States, 1795-1905*. H. Holt and Company.

Toivanen, Otto, & Väänänen, Lotta. 2016. Education and invention. *Review of Economics and Statistics*, **98**(2), 382–396.

United Nations, Dept of Social Affairs. 1953. *The Determinants and Consequences of Population Trends: A Summary of the Findings of Studies on the Relationships Between Population Changes and Economic and Social Conditions*. Vol. 17. UN.